arXiv:1510.05757v1 [math.GT] 20 Oct 2015

# Abelianizing $\mathrm{SL}_2\,\mathbb{R}$ local systems on compact surfaces

Aaron Fenyes

**Abstract**

The *abelianization* process introduced by Gaiotto, Hollands, Moore, and Neitzke turns $\mathrm{SL}_K\,\mathbb{C}$ local systems on a punctured surface into $\mathbb{C}^\times$ local systems, giving coordinates on the decorated $\mathrm{SL}_K\,\mathbb{C}$ character variety that are known to match Fock and Goncharov's cluster coordinates in the $\mathrm{SL}_2\,\mathbb{C}$ case. This paper extends abelianization to $\mathrm{SL}_2\,\mathbb{R}$ local systems on compact surfaces, using tools from dynamics to overcome the technical challenges that arise in the compact case. The approach taken here seems to complement the one recently used by Bonahon and Dreyer to arrive at a similar construction in a different geometric setting.

# 1 Introduction

## 1.1 Context

### 1.1.1 A sketch of the $\mathrm{SL}_2\,\mathbb{C}$ Hitchin systems

The Hitchin systems are a family of complex integrable systems, of interest both to representation theorists and to physicists studying supersymmetric gauge theories. The ones that will be most directly relevant to us describe "twisted" flat $\mathrm{SL}_2\,\mathbb{C}$ vector bundles on Riemann surfaces. I'll give only a quick sketch of how they look; a more detailed picture can be found in [1].

The *twisted* $\mathrm{SL}_2\,\mathbb{C}$ *character variety* of a compact Riemann surface $C$, which I'll denote $\mathcal{M}C$, is the space of irreducible flat $\mathrm{SL}_2\,\mathbb{C}$ vector bundles on the projective tangent bundle $\mathbf{P}TC$. The topological type of a bundle in $\mathcal{M}C$ is determined by its holonomy around the fiber of $\mathbf{P}TC$, which turns out to always be a square root of unity times the identity [2]. Hence, classifying bundles by topological type splits the character variety into two pieces, $\mathcal{M}_1C$ and $\mathcal{M}_{-1}C$, labeled by the square roots of unity.

Through a complicated chain of events, picking an element of $\mathcal{M}C$ stiffens the complex structure of $C$ to a *half-translation* structure (see Section 3.2 for a definition). This gives a map from $\mathcal{M}C$ to the space of half-translation structures on $C$, which is holomorphic with respect to a certain complex structure on $\mathcal{M}C$.[1] With the right holomorphic symplectic structure, this map makes $\mathcal{M}C$ into a complex algebraic integrable system. In the context of $\mathrm{SL}_2\,\mathbb{C}$ Hitchin

[1] In fact, there's a whole $\mathbb{C}^\times$-worth of such maps, each holomorphic with respect to its own associated complex structure. We just need one of them, say the one at $1 \in \mathbb{C}^\times$.

systems, the space of half-translation structures on $C$ is often called the *Hitchin base*.

If $C$ is a punctured Riemann surface, we can get a similar story by considering the *decorated character variety*: the space of irreducible flat $\mathrm{SL}_2\,\mathbb{C}$ vector bundles on $\mathbf{P}TC$ decorated with boundary data at the punctures. The choice of decoration at each puncture is discrete, so the decorated character variety is a finite cover of the plain one. A puncture with the simplest kind of decoration, a projectively flat section on a neighborhood, is called a *regular singularity*. Punctures with more complicated decorations are called *irregular singularities*. Picking an element of the decorated character variety gives $C$ a half-translation structure whose boundary conditions at the punctures are determined by the decorations. Each regular singularity, for instance, gets a neighborhood isometric to a half-infinite cylinder.

### 1.1.2 Abelianization and spectral coordinates

For a punctured Riemann surface, picking an element of the Hitchin base gives a remarkable coordinate system on a dense open subset of $\mathcal{M}_{-1}C$. This *spectral coordinate system* can be built using the general methods of Fock and Goncharov [3], but it gets its name from a more specialized construction discovered by Gaiotto, Hollands, Moore, and Neitzke. In this construction, the spectral coordinates arise as a byproduct of an *abelianization* process that turns twisted flat $\mathrm{SL}_2\,\mathbb{C}$ bundles on $C$ into twisted flat $\mathbb{C}^\times$ bundles on a branched double cover $\Sigma$ of $C$ [4, §4][5, §10]. This double cover is the *translation double cover* given by the half-translation structure on $C$, which is what we picked when we picked an element of the Hitchin base.[2] The abelianization process is reversible, and its codomain, the twisted $\mathbb{C}^\times$ character variety of $\Sigma$, is shaped like an algebraic torus $(\mathbb{C}^\times)^{2\,\mathrm{genus}\,\Sigma}$. Thus, abelianization can be seen as a torus-valued coordinate chart on $\mathcal{M}_{-1}C$.

Abelianization, in one sentence, proceeds by lifting a twisted flat $\mathrm{SL}_2\,\mathbb{C}$ bundle on $C$ to the translation double cover, cutting it along lines determined by the translation structure, and regluing it with a shear that flattens two previously non-flat (and in fact non-continuous) sub-bundles, splitting it into a direct sum. A slightly more detailed description can be found in Section 6.1.

## 1.2 Results

### 1.2.1 Abelianization without punctures

In the original account of abelianization, the punctures in $C$ played an important technical role, but Gaiotto, Moore, and Neitzke noted that their construction ought to work even when $C$ is compact [5, open problem 7]. The main result of this paper is to confirm that abelianization can be carried out on a compact surface, at least for twisted flat $\mathrm{SL}_2\,\mathbb{R}$ bundles. From a naive point of view,

[2] In [5], the authors go the other way, describing the *nonabelianization* process that turns a flat $\mathbb{C}^\times$ bundle into a twisted flat $\mathrm{SL}_2\,\mathbb{C}$ bundle with specified boundary data.

the process works just as it did in the punctured case. The catch is that the cutting and shearing is done along lines that fill $\Sigma$ densely, which is technically challenging to make sense of. The bulk of the paper will be spent building or acquiring the necessary machinery, and laying out the conditions under which it can be used successfully.

The main result, with all those subtleties accounted for, is summarized below. Its statement is very condensed, making it look intimidatingly remote, but the trip from here to there will hopefully feel more like a long hike up a gentle slope than a short climb up a sheer cliff. Words and ideas that won't be introduced until later along the route are tagged with references to the relevant sections. You can get an idea of what those sections are about from the table of contents in Section 1.3.

**Theorem 1.2.A** (Sections 6.2 – 6.3)**.** *Let $\Sigma$ be a compact translation surface (3.2.1) with generic dynamics (6.2), and let $\mathfrak{B}$ be its finite set of singularities. Let $\overleftrightarrow{\Sigma}$ be the associated divided surface (3), whose category of $\mathrm{SL}_2\mathbb{R}$ local systems is equivalent to the category of $\mathrm{SL}_2\mathbb{R}$ local systems on $\Sigma \smallsetminus \mathfrak{B}$ (Theorem 3.4.E).*

*Given a uniform (5) $\mathrm{SL}_2\mathbb{R}$ local system $\mathcal{E}$ on $\overleftrightarrow{\Sigma}$, we can find a new $\mathrm{SL}_2\mathbb{R}$ local system $\mathcal{F}$ and a stalkwise isomorphism $\Upsilon\colon \mathcal{E} \to \mathcal{F}$, supported on a dense subspace of $\overleftrightarrow{\Sigma}$, with the following properties:*

- *The deviation (2.3 – 2.4) of $\Upsilon$ behaves like the deviation of an abelianized local system from its original would behave on a punctured surface (6.1).*
- *The local system $\mathcal{F}$ splits into a direct sum of $\mathbb{R}^\times$ local systems.*

*Proof.* Sections 7 – 8. □

The most striking new feature of this result is the appearance of *uniformity*, a dynamical condition that plays no noticeable role in abelianization on punctured surfaces. The need for "generic dynamics" on $\Sigma$ is also new: the corresponding requirement in the punctured case is purely topological. As we'll see, abelianization on a compact surface involves dynamics in an essential way. This raises some interesting questions. On a punctured surface, for instance, deciding which local systems can be abelianized is pretty straightforward. The set of abelianizable local systems is open and dense, with a simple shape described by the opening sentence of Section 6.1.3. On a compact surface, deciding which local systems are uniform is much trickier. The set of uniform local systems is open, but not expected to be dense, and the descriptions I've seen of it suggest that its shape could be complicated. Section 5.5 and its main reference [7] say more about this.

### 1.2.2 Related work

In this paper, abelianization is done with a gadget similar to a transverse cocycle, which we'll call a *deviation*. The deviation that abelianizes a given $\mathrm{SL}_2\mathbb{R}$ local system should contain essentially the same information as the *shearing cycle* recently introduced by Bonahon and Dreyer for $\mathrm{SL}_n\mathbb{R}$ Hitchin characters [6].

The shearing cycle, however, lives in a different geometric setting: Bonahon and Dreyer work with Anosov representations on hyperbolic surfaces, while we'll be working with uniform local systems on translation surfaces.

Preliminary attempts to fill in the dictionary between these two approaches have been an interesting exercise in analogies between hyperbolic and flat geometry. A few points of contact will be highlighted as we encounter them.

### 1.2.3 Why not $\mathrm{SL}_2\,\mathbb{C}$?

I expect the results of this paper to generalize from $\mathrm{SL}_2\,\mathbb{R}$ to $\mathrm{SL}_2\,\mathbb{C}$ local systems. Outside of Section 5.5, where there may be some subtlety in saying what it means for an $\mathrm{SL}_2\,\mathbb{C}$ cocycle to be eventually positive, the generalization should amount to little more than declaring all the vector spaces to be complex. In our arguments about dynamical cocycles, however, we'll use results found in [7], [8], and [9], which only discuss the $\mathrm{SL}_2\,\mathbb{R}$ case. Until it can be verified that these results apply to $\mathrm{SL}_2\,\mathbb{C}$ cocycles, the results of this paper can only be stated with confidence for $\mathrm{SL}_2\,\mathbb{R}$ local systems.

## 1.3 Contents

***Introduction***

**Section 1** Context, summary of results, and various administrative things, including notation that will be used throughout the paper.

***Tools***

**Section 2** Abstracting from the idea of deforming a flat connection on a smooth bundle, we get a more general way of deforming local systems, called *warping*.

**Section 3** After a brief introduction to translation surfaces, we describe a way to enlarge a singular translation surface by splitting its critical leaves. We show that the resulting *divided surface* has useful topological and dynamical properties, and that it resembles the original surface both dynamically and in terms of its local systems.

**Section 4** We show how the warping process from Section 2 can be used to make sense of the idea of cutting and gluing a local sytem along the critical leaves of a divided surface, even when the critical leaves fill the surface densely.

**Section 5** We take a well-studied condition on dynamical cocycles, called *uniformity*, and reinterpret it as a condition on local systems on compact translation surfaces.

***Abelianization***

**Section 6** We review how abelianization works for $\mathrm{SL}_2\,\mathbb{C}$ local systems on a translation surface with punctures, point out the obstacles to carrying it out a compact translation surface, and describe how these obstacles can be overcome using the machinery from the previous sections. Section 6.3 contains the main product of the paper: instructions for abelianizing an $\mathrm{SL}_2\,\mathbb{R}$ local system on a compact translation surface, which are guaranteed to work under the conditions laid out in Section 6.2.

**Section 7** We show that abelianization, as defined by the instructions in Section 6.3, produces a well-defined local system, assuming the conditions from Section 6.2.

**Section 8** We show that the abelianized local system splits into a direct sum of $\mathbb{R}^\times$ local systems, assuming the conditions from Section 6.2.

**Section 9** A non-rigorous sample computation that uses abelianization to find holomorphic coordinates on the $\mathrm{SL}_2\,\mathbb{C}$ character variety of the punctured torus.

***Future directions***

**Section 10** We discuss a few of the interesting features that abelianization on a compact surface is expected to have, now that we know it can be done.

***Appendices***

**Appendix A** A pair of small technical lemmas for Section 2.

**Appendix B** A formalism for dynamical systems described by relations, used throughout the paper.

**Appendix C** Results on infinite ordered products, used heavily in Sections 4, 7, and 8.

**Appendix D** Linear algebra facts about the Euclidean plane, used in Section 8.

**Appendix E** A list of standard puncture shapes for translation surfaces and an explanation of where they come from, included to clarify the review in Section 6.

## 1.4 Setup

### 1.4.1 Running notation

The terminology of this section hasn't been introduced yet, but will be familiar to readers familiar with translation surfaces. If you'd like to become familiar with translation surfaces, skip ahead to Section 3.2.

From now on, $\Sigma$ will be a compact translation surface, $\mathfrak{B}$ its set of singularities, and $\mathfrak{W}$ the union of its critical leaves. Within $\mathfrak{W}$, let $\mathfrak{W}^+$ and $\mathfrak{W}^-$ be the unions of the backward- and forward-critical leaves, respectively. Saying that $\Sigma$ has no saddle connections is the same as saying that $\mathfrak{W}^+$ and $\mathfrak{W}^-$ are disjoint. The $\pm$ labeling is meant to evoke the fact that, in the absence of saddle connections, the vertical flow is well-defined on $\mathfrak{W}^+$ for all positive times, and on $\mathfrak{W}^-$ for all negative times.

I should stress that $\Sigma$ doesn't need to be the translation double cover of a half-translation surface, and there are nice examples of abelianization where it isn't. One of these is discussed in Section 9. There are many special properties abelianization is expected to gain when $\Sigma$ is a translation double cover, but our discussion of them will be limited to the speculative Section 10.

### 1.4.2 Index of symbols

Symbols can be hard to look up, so here's a list of unusual symbols that appear frequently in this paper, with references to the sections where they're defined. The first three are introduced in this paper, and the fourth is common, but not universal, in analysis.

- $\overleftrightarrow{\square}$ Divided interval or surface (Sections 3.3.1 and 3.4.1).
- $\overleftrightarrow{\square}$ Fractured interval or surface (same sections).
- $\overleftrightarrow{\square}$ Intersection with fractured interval or surface (Section 3.4.2).
- $\square \lesssim \square$ Bounded by a constant multiple (Section 7.1).

## 1.5 Acknowledgements

This paper would not have been started or finished without the tireless guidance of Andy Neitzke. I'm very grateful to Jen Berg for pointing out the argument used to prove Proposition 2.2.A, and to Sona Akopian for pointing out a key step in the proof of Theorem 3.3.K (I am, of course, responsible for any mistakes in these arguments). I've also enjoyed the benefit of conversations, some short and some long, with David Ben-Zvi, Francis Bonahon, Luis Duque, Richard Hughes, Tim Magee, Tom Mainiero, Taylor McAdam, Amir Mohammadi, Javier Morales, and Max Riestenberg, as well as diction brainstorming with Eliana Fenyes. This research was supported in part by NSF grants 1148490 and 1160461.

# 2 Warping local systems

## 2.1 Overview

In Section 1.1.2, I described abelianization as a cutting, shifting, and regluing process acting on flat bundles. On the punctured surfaces where abelianization was originally studied, this point of view works great, because the cutting and gluing only has to be done along a few isolated lines. On a compact surface,

however, the cutting and gluing happens along lines that fill the surface densely, and keeping track of what it does to the topology of the bundle becomes a serious hassle. To avoid it, we'll switch from studying flat bundles to studying locally constant sheaves, which have no explicit total space to worry about. Locally constant sheaves are also less sensitive to the topology of the space they live over, a feature we'll take great advantage of later.

In this section, we'll develop a general tool for deforming locally constant sheaves, called *warping*. We'll see in Section 4 that warping includes our dense cutting and gluing as a special case.

## 2.2 Conventions for local systems

Let's say $G$ is a $k$-linear group—a subgroup of the automorphism group of some finite-dimensional vector space $R_G$ over the field $k$. The object $R_G$ and the morphisms $G$ form a subcategory of $\mathbf{Vect}_k$, but it's a very small and lonely one. To broaden our horizons, let's define $\mathcal{C}_G$ to be the smallest subcategory of $\mathbf{Vect}_k$ which contains $R_G$ and $G$ and has all limits and filtered colimits. For our purposes, a $G$ *local system* will be a locally constant sheaf into $\mathcal{C}_G$ whose stalks are all isomorphic to $R_G$. Until Section 5, it won't matter much what $G$ is, so we'll often just talk about local systems in general.

The subcategory $\mathcal{C}_G$ is a nice target category for sheaves, because, like $\mathbf{Vect}_k$, it's a type of algebraic structure [10, Tag 007L]. Throughout this article, "sheaf" will mean a sheaf whose target category is a type of algebraic structure. With this said, we can define a constant sheaf to be a sheaf of locally constant functions—that is, functions constant on every connected component of their domain. In a locally connected space, every connected component of an open set is open, so the constant sheaf with value $A$ is characterized by the property that it sends every connected, non-empty open set to $A$, and every inclusion of such sets to $1_A$.

If $\mathcal{F}$ is a constant sheaf on a locally connected space, the stalk restriction morphism $\mathcal{F}_{x\in X} \colon \mathcal{F}_x \leftarrow \mathcal{F}_X$ is an isomorphism for every $x \in X$. In our context, the converse is true as well:

**Proposition 2.2.A.** *Suppose $\mathcal{F}$ is a sheaf on a locally connected space $X$. If the stalk restriction $\mathcal{F}_{x\in X}$ is an isomorphism for every $x \in X$, then $\mathcal{F}$ is isomorphic to a constant sheaf.*

*Proof.* Let $\bar{\mathcal{F}}$ be the constant sheaf with value $\mathcal{F}_X$. For each $U \subset X$, the restriction $\mathcal{F}_{U\subset X}$ gives a morphism $\bar{\mathcal{F}}_U \to \mathcal{F}_U$, and these morphisms fit together into a natural transformation from $\bar{\mathcal{F}}$ to $\mathcal{F}$. This natural transformation induces an isomorphism on every stalk, so it's an isomorphism of the underlying sheaves of sets, and therefore an isomorphism of sheaves of algebraic structures. (Many thanks to Jen Berg for pointing out this argument.) □

To save ink, let's say an open subset of a space is *simple* with respect to a sheaf if the restriction of the sheaf to the subset is isomorphic to a constant sheaf.

Notice that the value of a $G$ local system on a simple open set is isomorphic to $R_G$.

We'll frequently and without fanfare make use of the fact that a sheaf defined on a basis for a topological space extends uniquely (up to canonical isomorphism) to a sheaf on the full poset of open sets [10, Tag 009H, Lemma 9].

## 2.3 Deviations of flat connections

Say $E \to X$ is a smooth bundle, $A$ and $A'$ are flat connections on $E$, and $\mathcal{E}$ and $\mathcal{E}'$ are the corresponding sheaves of flat sections. The stalk $\mathcal{E}_x$ is the space of germs of flat sections of $A$ at $x$. Sending each germ in $\mathcal{E}_x$ to the unique germ in $\mathcal{E}'_x$ that has the same value at $x$ gives an isomorphism $\Upsilon_x \colon \mathcal{E}_x \to \mathcal{E}'_x$.

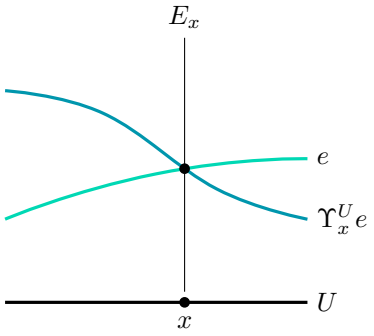


If $U \subset X$ is simple with respect to both $\mathcal{E}$ and $\mathcal{E}'$, we can visualize $\Upsilon_x$ by its action $\Upsilon_x^U \colon \mathcal{E}_U \to \mathcal{E}'_U$ on flat sections over $U$, as shown above. Then, for any $x, y \in U$, we can define an automorphism $\upsilon_{yx}^U$ of $\mathcal{E}_U$ that tells us how parallel transport along $A'$ deviates from parallel transport along $A$:

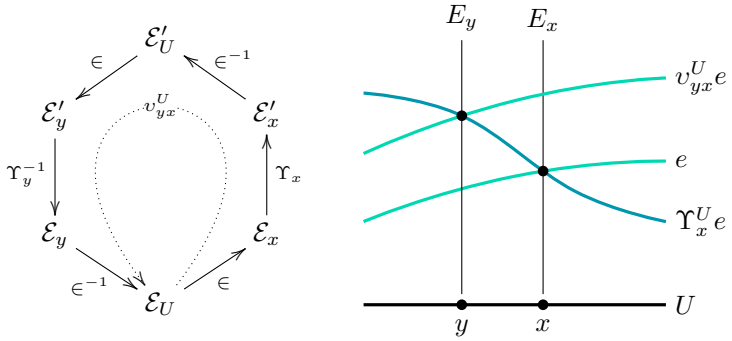


This automorphism is characterized by the property that $\Upsilon_y^U \upsilon_{yx}^U = \Upsilon_x^U$. If $V \subset U$ is a simple neighborhood of $x$ and $y$, the automorphisms $\upsilon_{yx}^U$ and $\upsilon_{yx}^V$ commute with the restriction map $\mathcal{E}_{V \subset U}$, so all these automorphisms fit together into a natural automorphism $\upsilon_{yx}$ of the functor we get by restricting $\mathcal{E}$ to the poset of simple neighborhoods of $x$ and $y$. Restricting $\mathcal{E}$ further to the simple

neighborhoods of three points $x$, $y$, and $z$, we can observe that $\upsilon_{zy}\upsilon_{yx} = \upsilon_{zx}$. Collectively, the natural automorphisms $\{\upsilon_{yx}\}_{x,y\in X}$ might be called the *deviation* of $A'$ from $A$.

## 2.4 Deviations of locally constant sheaves

More generally, say $X$ is just a locally connected space, and $\mathcal{F}$ and $\mathcal{F}'$ are locally constant sheaves on $X$. Given a stalkwise isomorphism $\Upsilon_x \colon \mathcal{F}_x \to \mathcal{F}'_x$, the construction of the natural automorphisms $\upsilon_{yx}$ goes through exactly as before. Since our choice of stalkwise isomorphism was not necessarily canonical, we should really talk about "the deviation of $\Upsilon$" instead of "the deviation of $\mathcal{F}'$ from $\mathcal{F}$."

Now, consider three locally constant sheaves $\mathcal{G}$, $\mathcal{F}$, and $\mathcal{F}'$ over $X$. If two stalkwise isomorphisms $\Phi \colon \mathcal{G} \to \mathcal{F}$ and $\Psi \colon \mathcal{G} \to \mathcal{F}'$ have the same deviation, $\upsilon$, how similar must they be? As it turns out, they're as good as identical: there's a unique natural isomorphism $T \colon \mathcal{F} \to \mathcal{F}'$ such that $\Psi_x = T_x \Phi_x$ for all $x \in X$. Here's why.

If we can find a natural isomorphism like this, it's clearly unique, because a natural transformation of sheaves is completely described by its action on stalks. Now, let's find one. Choose a basis $\mathcal{B}$ for the topology of $X$ consisting of sets which are simple with respect to all three sheaves. (We can do this because the sheaves are locally constant, and $X$ is locally conected.) Look at any basis element $U \in \mathcal{B}$. For any $x \in U$, define a morphism $T^U_x \colon \mathcal{F}_U \to \mathcal{F}'_U$ by

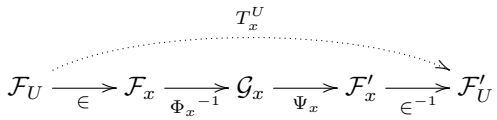


What if we had chosen another point $y \in U$ instead? Let $\upsilon$ be the shared deviation of $\Phi$ and $\Psi$, and consider the diagram

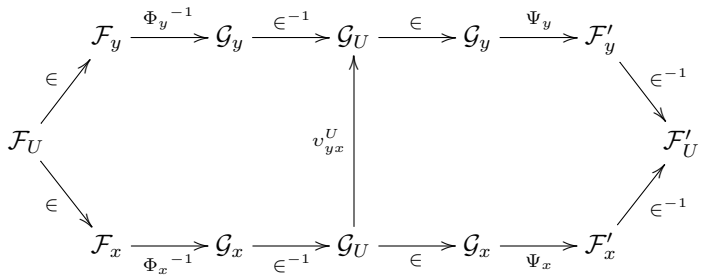


Notice that the bottom path is $T^U_x$, and the top path is $T^U_y$. The left and right chambers both commute, because $\upsilon$ is the deviation of both $\Phi$ and $\Psi$, so $T^U_x = T^U_y$ for all $x, y \in U$. Thus, we really have just one morphism $T^U \colon \mathcal{F}_U \to \mathcal{F}'_U$, which can be written in terms of any point in $U$.

For any other basis element $V \subset U$, writing $T^V$ and $T^U$ in terms of the same point $x \in V$ makes it easy to check that the square

$$\begin{array}{ccc} \mathcal{F}_U & \xrightarrow{T^U} & \mathcal{F}'_U \\ \subset\downarrow & & \downarrow\subset \\ \mathcal{F}_V & \xrightarrow[T^V]{} & \mathcal{F}'_V \end{array}$$

commutes. Since we've been working with arbitrary basis elements, we now see that the morphisms $\{T^U\}_{U \in \mathcal{B}}$ fit together into a natural transformation $T\colon \mathcal{F} \to \mathcal{F}'$, and it's clear by construction that $\Psi_x = T_x \Phi_x$ for all $x \in X$.

## 2.5 Warping locally constant sheaves

We just saw that, under favorable conditions, a stalkwise isomorphism is determined up to canonical isomorphism by its deviation. Let's see if we can go the other way and produce a stalkwise isomorphism with a specified deviation. First, we have to say what it means to specify a deviation.

Suppose $\mathcal{F}$ is a locally constant sheaf on a locally connected space $X$, $D$ is a dense subset of $X$, and $\mathcal{B}$ is a basis for the topology of $X$ consisting of $\mathcal{F}$-simple sets. To specify a *deviation* from $\mathcal{F}$ with support $D$, defined over the basis $\mathcal{B}$, we give for each pair of points $x, y \in D$ and each neighborhood $U \in \mathcal{B}$ of $x$ and $y$ an automorphism $\upsilon^U_{yx}$ of $\mathcal{F}_U$. These automorphisms have to fit together as follows:

- If $V \subset U$ is a basis element containing $x$ and $y$, the automorphisms $\upsilon^U_{yx}$ and $\upsilon^V_{yx}$ commute with the restriction morphism $\mathcal{F}_{V \subset U}$.
- For any three points $x, y, z \in D$, we have $\upsilon^U_{zy} \upsilon^U_{yx} = \upsilon^U_{zx}$.

The first condition just says that $\upsilon_{yx}$ is a natural automorphism of the restriction of $\mathcal{F}$ to the poset of basis elements containing both $x$ and $y$.

It turns out that, given a deviation $\upsilon$ from $\mathcal{F}$, we can always produce a locally constant sheaf $\mathcal{F}'$ and a stalkwise isomorphism $\Upsilon\colon \mathcal{F} \to \mathcal{F}'$, supported on $D$, whose deviation is $\upsilon$. We'll call this process *warping*. Here's how it's done. For each $U \in \mathcal{B}$, pick a point $x_U \in U \cap D$. Define $\mathcal{F}'_U$ to be the same as $\mathcal{F}_U$, but with the warped restriction morphism

$$\mathcal{F}'_{V \subset U} = \mathcal{F}_{V \subset U}\ \upsilon^U_{x_V x_U}$$

for each basis element $V \subset U$. Because every basis element is $\mathcal{F}$-simple, $\mathcal{F}_{V \subset U}$ is an isomorphism, so $\mathcal{F}'_{V \subset U}$ is an isomorphism too. It follows that the stalk restriction $\mathcal{F}'_{x \in U}$ is an isomorphism for any $x \in U$, so every basis element is $\mathcal{F}'$-simple. The stalkwise isomorphism $\Upsilon\colon \mathcal{F} \to \mathcal{F}'$ is given by

$$\Upsilon_x = \mathcal{F}'_{x \in U}\ \upsilon^U_{x_U x}\ {\mathcal{F}_{x \in U}}^{-1}$$

for any basis element $U$ containing $x \in D$.

There are three claims implicit in the description of $\mathcal{F} \xrightarrow{\Upsilon} \mathcal{F}'$ above:

- $\mathcal{F}'$ is a locally constant sheaf.
- The definition of $\Upsilon_x$ doesn't depend on our choice of neighborhood $U$.
- The deviation of $\mathcal{F}'$ from $\mathcal{F}$ is $\upsilon$.

Let's check these claims.

**$\mathcal{F}'$ is a locally constant sheaf** The functoriality of $\mathcal{F}'$ follows easily from the fact that $\upsilon$ is a deviation. To verify that $\mathcal{F}'$ is a sheaf, pick any element $U$ of the basis $\mathcal{B}$. Suppose that for each basis element $V \subset U$, we have an element $s_V$ of $\mathcal{F}'_V$, and these elements commute with the restriction morphisms of $\mathcal{F}'$. We need to find an element $s$ of $\mathcal{F}'_U$ that restricts to $s_V$ on every basis element $V \subset U$. Since each of the restrictions $\mathcal{F}'_{V \subset U}$ is an isomorphism, there can only be one element like this, and it will exist if and only if the elements ${\mathcal{F}'_{V \subset U}}^{-1} s_V$ match for all the basis elements $V \subset U$.

For two basis elements $W \subset V$ contained in $U$, we have

$$\begin{aligned} {\mathcal{F}'_{W \subset U}}^{-1} s_W &= {\mathcal{F}'_{W \subset U}}^{-1} \mathcal{F}'_{W \subset V} s_V \\ &= {\mathcal{F}'_{V \subset U}}^{-1} s_V, \end{aligned}$$

so $W$ and $V$ give the same element. Since any two overlapping basis elements contain another basis element in their intersection, it follows that any two overlapping basis elements give the same element as well. Because $U$ is connected, any two basis elements can be linked by a finite sequence of overlapping basis elements (Appendix A.1). Therefore, all the basis elements contained in $U$ give the same element. This is the element $s$ we were looking for, completing our proof that $\mathcal{F}'$ is a sheaf.

To see that $\mathcal{F}'$ is locally constant, first recall that the elements of $\mathcal{B}$ are simple, so the restriction arrows of $\mathcal{F}$ over $\mathcal{B}$ are isomorphisms. Thus, the restriction arrows of $\mathcal{F}'$ are isomorphisms as well. Now, pick any point $x \in X$, not necessarily in $D$. Applying $\mathcal{F}'$ to the poset of basis elements containing $x$ yields a downward-directed diagram whose arrows are all isomorphisms. The defining arrows from a diagram like this to its colimit are always isomorphisms (Appendix A.2). In other words, the stalk restriction $\mathcal{F}'_{x \in U}$ is an isomorphism for every $U \in \mathcal{B}$ containing $x$. Since $x$ was an arbitrary point in $X$, it follows by Proposition 2.2.A that $\mathcal{F}'$ is locally constant.

**$\Upsilon_x$ is well-defined** To see that the definition of $\Upsilon_x$ doesn't depend on our choice of neighborhood, first observe that for two basis elements $V \subset U$ con-

taining $x$,

$$\begin{aligned}(\mathcal{F}'_{x\in U})\ \upsilon^U_{x_U x}\ \mathcal{F}_{x\in U}{}^{-1} &= (\mathcal{F}'_{x\in V}\ [\mathcal{F}'_{V\subset U}])\ \upsilon^U_{x_U x}\ \mathcal{F}_{x\in U}{}^{-1}\\ &= \mathcal{F}'_{x\in V}[\mathcal{F}_{V\subset U}\ (\upsilon^U_{x_V x_U}]\ \upsilon^U_{x_U x})\ \mathcal{F}_{x\in U}{}^{-1}\\ &= \mathcal{F}'_{x\in V}\ [\mathcal{F}_{V\subset U}\ (\upsilon^U_{x_V x})]\ \mathcal{F}_{x\in U}{}^{-1}\\ &= \mathcal{F}'_{x\in V}\ [\upsilon^V_{x_V x}\ (\mathcal{F}_{V\subset U}]\ \mathcal{F}_{x\in U}{}^{-1})\\ &= \mathcal{F}'_{x\in V}\ \upsilon^V_{x_V x}\ (\mathcal{F}_{x\in V}{}^{-1}),\end{aligned}$$

so $V$ and $U$ give the same isomorphism. Since any two basis elements containing $x$ contain another basis element in their intersection, it follows that every basis element containing $x$ gives the same isomorphism.

**The deviation of $\mathcal{F}'$ from $\mathcal{F}$ is $\upsilon$** Finally, let $\delta$ be the deviation of $\Upsilon$. It's easy to calculate $\delta^U_{yx}$ by defining $\Upsilon_x$ and $\Upsilon_y$ in terms of $U$:

$$\begin{aligned}\delta^U_{yx} &= \mathcal{F}_{y\in U}{}^{-1}\ (\Upsilon_y^{-1})\ \mathcal{F}'_{y\in U}\ \mathcal{F}'_{x\in U}{}^{-1}\ (\Upsilon_x)\ \mathcal{F}_{x\in U}\\ &= \mathcal{F}_{y\in U}{}^{-1}\ (\mathcal{F}_{y\in U}\ \upsilon^U_{yx_U}\ \mathcal{F}'_{y\in U}{}^{-1})\ \mathcal{F}'_{y\in U}\ \mathcal{F}'_{x\in U}{}^{-1}\ (\mathcal{F}'_{x\in U}\ \upsilon^U_{x_U x}\ \mathcal{F}_{x\in U}{}^{-1})\ \mathcal{F}_{x\in U}\\ &= \upsilon^U_{yx_U}\ \upsilon^U_{x_U x}\\ &= \upsilon^U_{yx}.\end{aligned}$$

Thus, $\delta = \upsilon$, as claimed.

### 2.6 Warping local systems

Local systems are just a special kind of locally constant sheaves, so all the constructions of the previous sections can be applied to them. In this case, the definition of a deviation can be pared down a bit, because if $\mathcal{F}$ is a $G$ local system and $U$ is an $\mathcal{F}$-simple open set, an automorphism of $\mathcal{F}_U$ is just an element of $G$.

Warping a local system always produces another local system. To see why, take a $G$ local system $\mathcal{F}$ and warp it by some deviation. The warped sheaf $\mathcal{F}'$ is locally constant, and stalkwise isomorphic to $\mathcal{F}$ over the support of the deviation. Because the support is dense, it follows that every stalk of $\mathcal{F}'$ is isomorphic to a stalk of $\mathcal{F}$, and hence isomorphic to the vector space $R_G$ on which $G$ acts.

## 3 Dividing translation surfaces

### 3.1 Overview

For working out the technical details of abelianization, it will be useful to embed the surface $\Sigma \smallsetminus \mathfrak{B}$ in a larger space $\overset{\leftrightarrow}{\Sigma}$, called the *divided surface*, whose local systems are naturally in correspondence with the local systems on $\Sigma \smallsetminus \mathfrak{B}$. On the divided surface, we can stand infinitesimally close to any critical leaf, streamlining our discussion of the abelianization process in Sections 6 and 7.

Removing the critical leaves of $\Sigma$ from the divided surface yields a compact space $\overleftrightarrow{\Sigma}$, called the *fractured surface*, which can be metrized in a very natural way. Its metric properties will play a crucial role in Section 8, where we prove that abelianization does the job it's meant to do.

The divided surface is also a point of contact between abelianization and the similar construction introduced by Bonahon and Dreyer, which takes place on a hyperbolic surface with a maximal geodesic lamination rather than a flat surface with a translation structure [6]. Divided surfaces can be seen as intermediates between these two kinds of surfaces, as sketched in Section 3.6.

## 3.2 A review of translation and half-translation surfaces

### 3.2.1 Translation surfaces

A non-singular *translation surface* is a manifold whose charts are open subsets of $\mathbb{R}^2$ and whose transition maps are translations. Every translation surface comes with a bunch of geometric structures induced by the translation-invariant geometric structures on $\mathbb{R}^2$, which include:

- The flat metric.
- The four cardinal directions: up, down, right, and left.
- The vertical and horizontal foliations, whose leaves are vertical and horizontal lines. Both foliations can be oriented; we'll orient them upward and rightward, respectively.
- The vertical flow, which moves points upward at unit speed. On a surface which is non-compact, as most non-singular translation surfaces are, this flow might not be defined everywhere at all times. In general, the flow at a given time will be only a *bicontinuous relation*, rather than a homeomorphism (see Appendix B for details).

Whenever I refer to a foliation of a translation surface, I mean the vertical one, unless I say otherwise.

It's conventional, and convenient, to allow translation surfaces to have *conical singularities*, which look like this:

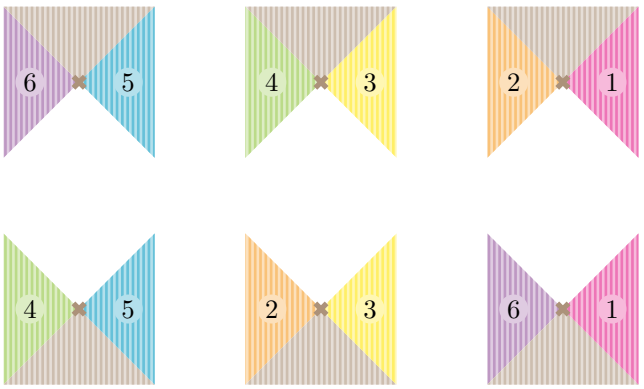


The like-numbered triangles are identified through translation. The triangles include the cross-marked center points, which get quotiented down to a single point by the identifications. In the case shown above, the total angle around the singularity is $6\pi$; in general, any even multiple of $\pi$ is possible. It's sometimes useful to mark a discrete set of ordinary points as "singularities" of cone angle $2\pi$.

A compact translation surface can have only finitely many singularities. Conformally, the vertical and horizontal foliations in the neighborhood of a singularity look like this:

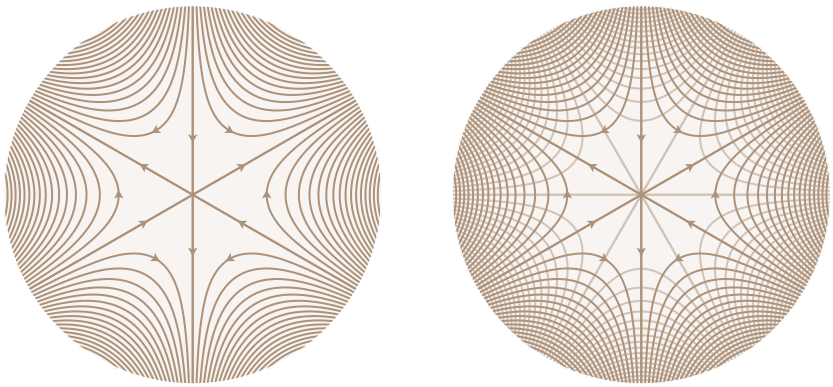

Vertical



Vertical and horizontal

The vertical leaves that dive into the singularity are called *forward-critical*, and the ones that shoot out of the singularity are called *backward-critical*. A leaf which is critical in both directions is called a *saddle connection*. Critical leaves are spaced evenly around the singularity at angles of $\pi$.

The vertical flow on a singular translation surface is only defined away from the singularities. It acts by bicontinuous relations, making points on the critical leaves disappear as they fall into the singularities. When restricted to the complement of the critical leaves, the vertical flow acts by homeomorphisms.

A translation surface is said to be *minimal* if all its vertical leaves are dense. We'll see in Section 3.2.2 that a non-critical leaf which is dense in one direction must be dense in both directions. Thus, on a minimal translation surface, every leaf that's not forward-critical is dense in the forward direction, and every leaf that's not backward-critical is dense in the backward direction. A translation surface with no saddle connections is automatically minimal [11, proof of Theorem 1.8].

Away from the singularities, the topology of a translation surface has a basis consisting of *flow boxes*: open rectangles with vertical and horizontal sides. A compact translation surface can be covered by a finite collection of flow boxes and singularity charts. The special class of *well-cut* flow boxes, defined in the next section, will play an important role in this paper.

Because translations form a normal subgroup of $\mathrm{Aff}\,\mathbb{R}^2$, a translation structure can be modified by composing an element of $\mathrm{GL}_2\,\mathbb{R}$ with all its charts. In particular, a translation structure can be rotated, tilting the vertical foliation. If you rotate a translation structure through a full circle, all but countably many of the structures you pass through will have no saddle connections [11, proof of Theorem 1.8]. Moreover, for all but a measure-zero subset of the minimal structures, the vertical flow will be uniquely ergodic on the complement of the critical leaves [11, Theorem 3.5]. Hence, an arbitrarily small rotation is all it takes to turn any translation structure into a translation structure with no saddle connections and a vertical flow which is uniquely ergodic on the complement of the critical leaves.

### 3.2.2 First return maps

Let's say a *horizontal segment* on a translation surface is a subset that looks like an open, closed, or half-open horizontal line segment in some chart. Pick a point on a horizontal segment, and watch it as it's carried upward by the vertical flow. On a compact translation surface, unless it falls into a singularity, the point will eventually return to the segment it started on. This is an immediate corollary of [11, Lemma 1.7], which for reference I'll restate here.

**Lemma 3.2.A.** *Let $Z$ be a closed horizontal segment on a compact translation surface, and let $p$ be one of its endpoints. Unless the vertical leaf through $p$ is forward-critical, the vertical flow will eventually carry $p$ back to $Z$.*

On any horizontal segment $Z$ in a compact translation surface, we can define a relation $\alpha$ that sends each point to the place where it first returns to $Z$ under the vertical flow. When fed a point that falls into a singularity before returning, $\alpha$ gives back nothing. We'll call $\alpha$ the *first return relation* on $Z$. The inverse relation $\alpha^{-1}$ sends each point to the place where it first returns to $Z$ under the backward vertical flow.

When restricted to the complement of the forward-critical leaves, $\alpha$ becomes a function, and is called the *first return map*. Similarly, $\alpha^{-1}$ becomes a function when restricted to the complement of the backward-critical leaves. On the complement of all the critical leaves, $\alpha$ and $\alpha^{-1}$ are inverse functions.

Because the vertical foliation only has a few kinds of local geometry, the first return relation only has a few kinds of local behavior. Under mild conditions, it belongs to the class of transformations called *interval exchanges*, which we'll hear more about in Section 3.3.2 [12, §3]. The conditions are the price we pay for defining first return relations, horizontal segments, and interval exchanges in a slightly non-standard way, whose advantages will become apparent in Section 3.5.2.

In our framework, the first return relation is only an interval exchange on a *well-cut* segment, which is a horizontal segment with the following properties:

- It looks like an open horizontal line segment bounded by critical leaves. (Both endpoints may lie on the same critical leaf.)
- The forward vertical flow drops each forward-critical boundary point into a singularity without carrying it through the segment. Similarly, the backward vertical flow drops each backward-critical boundary point into a singularity without carrying it through the segment.

When we construct the forward and backward first return relations on a horizontal segment $Z$, this condition prevents the vertical flow $\psi^t$ from breaking $\psi^t Z$ across the boundaries of $Z$ in a way inconsistent with our definition of an interval exchange. We'll call a flow box *well-cut* if every horizontal slice across it is a well-cut segment.

When critical leaves are plentiful, well-cut segments are easy to find.

**Proposition 3.2.B.** *If $Z$ is an open horizontal segment bounded by critical leaves, every non-critical point on $Z$ is contained in a well-cut subsegment of $Z$.*

*Proof.* Carry each forward-critical boundary point of $Z$ along the forward vertical flow, marking each place it passes through $Z$ before falling into a singularity. Do the same with each backward-critical boundary point, using the backward vertical flow. Removing the marked points breaks $Z$ into a finite collection of open horizontal segments, which are all well-cut. □

### 3.2.3 Half-translation surfaces

A half-translation surface is the same thing as a translation surface, but with transition maps composed of both translations and half-turn rotations. This expansion of the structure group is pretty small, so half-translation surfaces have almost as much structure as translation surfaces do:

- The flat metric remains.
- The vertical and horizontal foliations remain, but they're no longer canonically oriented. In fact, it's often not possible to give either foliation a consistent global orientation.
- As a result, it's often not possible to define a global vertical flow.

The group generated by translations and half-turns has the translation group as a normal subgroup, so each of its elements is either a translation or a translation followed by a half-turn. I'll call the latter a *flip*.

In a translation surface, we saw that the angle around a conical singularity could be any even multiple of $\pi$. In a half-translation surface, any multiple of $\pi$ is possible. Here's a pattern for a singularity with cone angle $3\pi$:

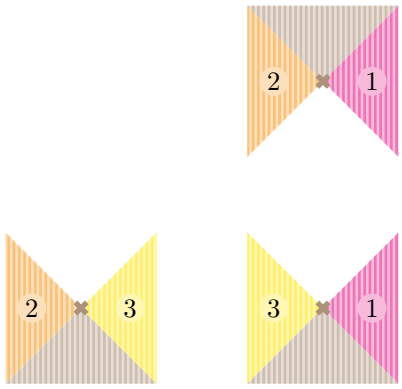


The triangles labeled 3 are identified by a flip. A singularity with cone angle $\pi$ can be constructed, informally, by using a flip to glue one of the notched square building blocks we've been using to itself:

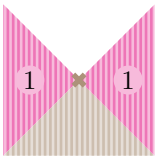


The vertical and horizontal foliations can't be oriented consistently around an odd singularity, but they can be oriented on a small enough neighborhood of any other point. In particular, on a half-translation surface, the definition of a flow box still makes sense, and the vertical and horizontal foliations can be oriented inside a flow box.

Every half-translation surface comes with a translation surface hovering over it as a branched double cover, with a projection map that preserves the half-translation structure. Away from the singularities, this *translation double cover* is built by making two copies of each flow box on the half-translation surface, one for each possible orientation of the vertical foliation. The transition maps of the half-translation surface induce transitions between the oriented copies in a natural way.

The translation double cover can be extended over the whole surface by duplicating the cut square pieces that make up the region around each singularity. In the picture below, duplicate pieces are labeled with the same numbers. One piece from each pair is drawn in dark ink, and the other in light ink.

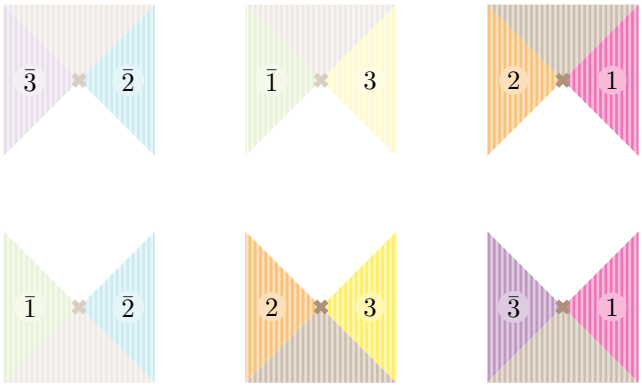


The new region is a double cover of the original, branched at the singularity. The branch point is a conical singularity with twice the angle of the original. The doubling process produces a branch point even when you start with a singularity of cone angle $2\pi$, so we've found a meaning for these removable singularities: they stand for genuine singularities in the translation double cover. At the same time, the doubling process turns singularities with cone angle $\pi$ into ones with cone angle $2\pi$, so removable singularities above mark genuine singularities below.

## 3.3 Dividing intervals

### 3.3.1 Construction of divided and fractured intervals

Dividing a translation surface is essentially a one-dimensional process, so let's start in the one-dimensional case. Let $I \subset \mathbb{R}$ be an open interval, and let $W$ be a subset of $I$ which is countable or smaller. To *divide* $I$ at $W$, first build a new set

$$\overleftrightarrow{I} = \{\overleftarrow{w}, \hat{w}, \vec{w}\}_{w \in W} \sqcup \{s\}_{s \in I \smallsetminus W}.$$

There's an obvious map $\pi \colon \overleftrightarrow{I} \to I$ which sends $\overleftarrow{w}$, $\hat{w}$, and $\vec{w}$ to $w$ and each point in $I \smallsetminus W$ to itself. Order $\overleftrightarrow{I}$ so that $\pi$ is order-preserving and $\overleftarrow{w} < \hat{w} < \vec{w}$. Give $\overleftrightarrow{I}$ the topology generated by all the non-empty intervals $(a, b)$ except the ones that look like $(\overleftarrow{w}, b)$ or $(a, \vec{w})$. This topology is coarser than the order topology, and non-Hausdorff: neither $\overleftarrow{w}$ nor $\vec{w}$ can be separated from $\hat{w}$ by open sets. The generating intervals described above form a basis for the topology, so I'll refer to them as "basis intervals."

As you might expect, $\pi$ turns out to be a quotient map. Going the other direction, let $\iota \colon I \to \overleftrightarrow{I}$ be the map that sends $w \in W$ to $\hat{w}$ and each point in $I \smallsetminus W$ to itself. Perhaps surprisingly, $\iota$ turns out to be an embedding. (Both of these claims will be proven in Section 3.3.3.) Define the *fractured interval* $\overleftrightarrow{I}$ to be the complement of $\iota W$ in the divided interval $\overleftrightarrow{I}$.

### 3.3.2 Examples from dynamics

Divided intervals arise naturally in the study of some one-dimensional dynamical systems. In this setting, the fractured interval often turns out to be a familiar coding of the system. The binary shift provides an excellent example.

Divided intervals arising from interval exchange transformations will play an important role in this paper. The corresponding fractured intervals, for minimal interval exchanges, have appeared in the literature as *Cantor minimal systems* [13].

In both examples, we'll describe the dynamics using a partial map; you'll probably be able to guess from context what that means. To be precise, it means a coinjective, bicontinuous relation, in the terminology of Appendix B.

**The binary shift** The *binary shift* is the partial map from $(0,1)$ to itself that sends $s$ to $\lfloor 2s \rfloor$, returning nothing if $\lfloor 2s \rfloor = 1$. Shifting a number in $(0,1)$ removes the first digit of its binary expansion. We can extend the binary shift to all reals by thinking of $\mathbb{R} \smallsetminus \mathbb{Z}$ as a union of copies of $(0,1)$. Shifting a number in $\mathbb{R} \smallsetminus \frac{1}{2}\mathbb{Z}$ shifts the fractional part of its binary expansion, leaving the integer part alone. Shifting a number in $\frac{1}{2}\mathbb{Z}$ returns nothing.

Applying the shift map over and over, let $W \subset \mathbb{R}$ be the set of points that eventually "fall into a break," reaching a point where the map returns nothing. This turns out to be the set of rationals whose denominators are powers of two.

Divide $\mathbb{R}$ at $W$. Let $C \subset \overleftrightarrow{\mathbb{R}}$ be the interval $[\vec{0}, \overleftarrow{1}]$, noting that $\pi C = [0,1]$, and let $\check{C} = C \cap \check{\mathbb{R}}$. We'll soon learn, from Corollary 3.3.I in Section 3.3.4, that $\check{C}$ is a Cantor set. In the meantime, we can see this directly by identifying $\check{C}$ with the space of one-sided binary sequences. A number $s \in (0,1) \smallsetminus W$ has a unique binary expansion, $\iota s$. A number $w \in W$ has two binary expansions: $\overleftarrow{w}$, the one ending in ones, and $\vec{w}$, the one ending in zeros. The quotient map $\pi \colon \check{C} \to [0,1]$ interprets each binary sequence as a number.

The *shift* on a space of one-sided sequences is the map that removes the first symbol of the sequence. The shift on $\check{C}$ and the binary shift on $(0,1) \smallsetminus W$ commute with the embedding $\iota \colon (0,1) \smallsetminus W \to \check{C}$. The shift on $\check{C}$ extends uniquely to a relation on $C$ that commutes with $\iota \colon (0,1) \to C$. This relation can be seen as a divided version of the binary shift. Notice that the binary shifts on $C$ and $[0,1]$ don't quite commute with the quotient map $\pi$, because when a point $w \in W$ falls into a break, its lifts $\overleftarrow{w}$ and $\vec{w}$ keep going.

**Interval exchanges** An *interval exchange transformation* is a partial map from $(0,1)$ to itself that works by splitting $(0,1)$ into finitely many open subintervals and shuffling them around:

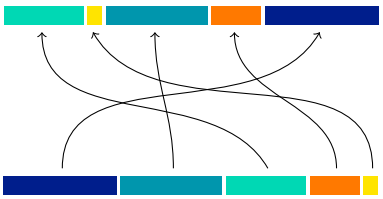

On the break points between the intervals, the map returns nothing.

An interval exchange, unlike the binary shift, is injective, so its inverse relation is also a partial map. In fact, its inverse is another interval exchange, which unshuffles the pieces of $(0, 1)$.

Just as we did with the binary shift, we can extend an interval exchange to all reals by thinking of $\mathbb{R} \setminus \mathbb{Z}$ as a union of copies of $(0, 1)$.

Pick an interval exchange, and let $W \subset \mathbb{R}$ be the set of points that eventually fall into a break under iteration of either this map or its inverse. Just as before, divide $\mathbb{R}$ at $W$. Let $C \subset \overleftrightarrow{\mathbb{R}}$ be the interval $[\vec{0}, \overleftarrow{1}]$, noting that $\pi C = [0, 1]$, and let $\dot{\dot{C}} = C \cap \dot{\dot{\mathbb{R}}}$. If $W$ is dense in $\mathbb{R}$, Corollary 3.3.I in Section 3.3.4 will tell us that $\dot{\dot{C}}$ is a Cantor set. In this case, $\dot{\dot{C}}$ can be seen as a subspace of the two-sided sequence space whose alphabet is the set of intervals shuffled by the interval exchange. The orbit of a point $s \in (0, 1) \setminus W$ under iteration of the interval exchange and its inverse is infinite in both directions, so we can get a two-sided sequence $\iota s$ by keeping track of which intervals the orbit passes through.

The orbit of a point $w \in W$ ends when it falls into a break, but it can be continued in two natural ways. One is to extend the interval exchange relation to a left-continuous map on $(0, 1]$, so each point travels with the points to the left of it, and the intervals being shuffled become closed on the right. The orbit of $w$ becomes infinite in both directions, and keeping track of which intervals it goes through yields a two-sided sequence $\overleftarrow{w}$. The other way to continue the orbit of $w$ is to extend the interval exchange relation to a right-continuous map on $[0, 1)$, giving a different two-sided sequence $\vec{w}$.

The *shift* on a space of two-sided sequences is the map that moves every symbol one step earlier. The shift map on $\dot{\dot{C}}$ and the interval exchange map on $(0, 1) \setminus W$ commute with the embedding $\iota\colon (0, 1) \setminus W \to \dot{\dot{C}}$. The shift on $\dot{\dot{C}}$ extends uniquely to a relation on $C$ that commutes with $\iota\colon (0, 1) \to C$. This relation can be seen as a divided version of the interval exchange. Notice that interval exchanges on $C$ and $[0, 1]$ don't quite commute with the quotient map $\pi$, because when a point $w \in W$ falls into a break, its lifts $\overleftarrow{w}$ and $\vec{w}$ keep going, following the left- and right-continuous extensions of the interval exchange.

### 3.3.3 Properties of divided intervals

For convenience, let $\overleftarrow{\iota}\colon I \to \overleftrightarrow{I}$ be the map that sends $w \in W$ to $\overleftarrow{w}$ and each point in $I \setminus W$ to itself. Define $\vec{\iota}$ similarly.

A basis interval has a leftmost element if and only if it looks like $(\hat{w}, b) = [\vec{w}, b)$, and a rightmost element if and only if it looks like $(a, \hat{w}) = (a, \overleftarrow{w}]$. It will often be useful to *trim* a basis interval by removing its leftmost and rightmost

elements, if they exist. The trimmed version of an interval $(a, b) \subset \overleftrightarrow{I}$, denoted $\mathrm{trim}(a, b)$, can be written explicitly as $(\vec{\iota}\pi a, \overleftarrow{\iota}\pi b)$. Notice that $\pi\,\mathrm{trim}(a, b) = \iota^{-1}(a, b) = (\pi a, \pi b)$ for any basis interval $(a, b)$, and that trimming a basis interval does not remove any points in the image of $\iota$. Conveniently, for any basis interval, $\pi^{-1}\iota^{-1}(a, b) = \mathrm{trim}(a, b)$.

With these tools in hand, let's prove the claims about $\pi$ and $\iota$ made in the previous section.

*Proof that $\pi$ is a quotient map.* To see that $\pi$ is continuous, observe that the preimage of $(a, b) \subset I$ under $\pi$ is the basis interval $(\vec{\iota}a, \overleftarrow{\iota}b)$.

To see that $\pi$ is a quotient map, pick any $S \subset I$ whose preimage under $\pi$ is open. We want to show $S$ is open. For any $s \in S$, the point $\iota s$ is in $\pi^{-1}S$, so there is a basis interval $H \subset \pi^{-1}S$ containing $\iota s$. Since $\mathrm{trim}\, H$ also contains $\iota s$, and $\pi$ sends trimmed basis intervals to open intervals, $\pi\,\mathrm{trim}\, H$ is an open subset of $S$ containing $s$. □

*Proof that $\iota$ is an embedding.* To see that $\iota$ is continuous, recall that $\iota^{-1}(a, b) = (\pi a, \pi b)$ for any basis interval $(a, b)$.

To see that $\iota$ is an embedding, observe that the image under $\iota$ of an interval $(a, b) \subset I$ is the intersection of $(\iota a, \iota b)$ with $\iota I$. □

The continuity of $\iota$ is a way of saying that passing from $I$ to $\overleftrightarrow{I}$ spreads out the points of $W$, but it doesn't spread them out too much. Here are two more reflections of this idea.

**Proposition 3.3.A.** *The embedding of $I$ in $\overleftrightarrow{I}$ is dense.*

*Proof.* It's enough to show that $\iota I$ intersects every basis interval. Suppose the basis interval $(a, b)$ doesn't intersect $\iota I$, so its preimage $(\pi a, \pi b)$ under $\iota$ is empty. Since $I$ is densely ordered, this means $\pi a = \pi b$, which is precluded by the rules defining basis intervals. □

**Proposition 3.3.B.** *The divided interval $\overleftrightarrow{I}$ is locally connected.*

*Proof.* It's enough to show that every basis interval is connected. Recall that basis intervals are non-empty by definition. Let's say the basis interval $(a, b)$ is disconnected by two open subsets $U$ and $V$. Since $\iota I$ is dense in $\overleftrightarrow{I}$, the preimages of $U$ and $V$ under $\iota$ are non-empty, so they disconnect the preimage $(\pi a, \pi b)$ of $(a, b)$. □

For our purposes, the most important feature of $\overleftrightarrow{I}$ is that its local systems are naturally in correspondence with the local systems on $I$. This idea can be stated more precisely as follows.

**Theorem 3.3.C.** *For any linear group $G$, the direct image functors $\pi_*$ and $\iota_*$ give an equivalence between the category of $G$ local systems on $\overleftrightarrow{I}$ and the category of $G$ local systems on $I$.*[3]

The reason $\overleftrightarrow{I}$ has no more local systems than $I$, despite having more open subsets, is that a local system on $\overleftrightarrow{I}$ is determined entirely by its values on trimmed intervals.

**Lemma 3.3.D.** *If $\mathcal{F}$ is a local system on $\overleftrightarrow{I}$, the restriction $\mathcal{F}_{\operatorname{trim} H \subset H}$ is an isomorphism for any basis interval $H$.*

*Proof.* If $H$ has neither a least element nor a greatest element, $\operatorname{trim} H = H$, so there's nothing to prove. Let's assume $H$ has a least element, but no greatest element; the remaining cases are essentially the same.

Since $\mathcal{F}$ is locally constant, we can pick a basis interval $A \subset H$ which contains the least element of $H$ and is small enough that $\mathcal{F}|_A$ is constant. Since $A$ and $\operatorname{trim} A$ are both connected, $\mathcal{F}_{\operatorname{trim} A \subset A}$ is an isomorphism. The diagram

$$\begin{array}{ccccc}
 & & \mathcal{F}_{\operatorname{trim} A} & \xleftarrow{\subset} & \mathcal{F}_{\operatorname{trim} H} \\
 & \overset{\subset^{-1}}{\swarrow} & \downarrow{\scriptstyle =} & & \downarrow{\scriptstyle =} \\
\mathcal{F}_A & \xrightarrow[\subset]{} & \mathcal{F}_{\operatorname{trim} A} & \xleftarrow[\subset]{} & \mathcal{F}_{\operatorname{trim} H}
\end{array}$$

commutes, so taking limits of its top and bottom rows gives a map $\mathcal{F}_{\operatorname{trim} H} \to \mathcal{F}_H$, which inverts $\mathcal{F}_{\operatorname{trim} H \subset H}$. □

*Proof of Theorem 3.3.C.* There's a canonical natural isomorphism between $\pi_*\iota_*$ and the identity functor, because $\pi\iota$ is the identity map from $I$ to itself. Now, all we need is a natural isomorphism between $\iota_*\pi_*$ and the identity.

Pick any local system $\mathcal{F}$ on $\overleftrightarrow{I}$. For each basis interval $H$, recall that $\pi^{-1}\iota^{-1}H = \operatorname{trim} H$, so

$$\begin{aligned}
(\iota_*\pi_*\mathcal{F})_H &= \mathcal{F}_{\pi^{-1}\iota^{-1}H} \\
&= \mathcal{F}_{\operatorname{trim} H}.
\end{aligned}$$

Thus, the restriction $\mathcal{F}_{\operatorname{trim} H \subset H}$ gives a morphism from $\mathcal{F}_H$ to $(\iota_*\pi_*\mathcal{F})_H$, and Lemma 3.3.D tells us this morphism is an isomorphism. For any basis interval $H' \subset H$, the diagram

$$\begin{array}{ccc}
\mathcal{F}_{H'} & \xleftarrow{\subset} & \mathcal{F}_H \\
\downarrow{\scriptstyle \subset} & & \downarrow{\scriptstyle \subset} \\
\mathcal{F}_{\operatorname{trim} H'} & \xleftarrow[\subset]{} & \mathcal{F}_{\operatorname{trim} H}
\end{array}$$

commutes because all the arrows are restrictions, so we've found a natural isomorphism from $\mathcal{F}$ to $\iota_*\pi_*\mathcal{F}$. □

[3] Though they're stated for categories of local systems, Theorem 3.3.C and Lemma 3.3.D hold for categories of locally constant sheaves into any fixed target category. The proofs are the same, keeping in mind our convention (from Section 2.2) that the target category is a type of algebraic structure.

### 3.3.4 Properties of fractured intervals

One nice feature of $\overleftrightarrow{I}$ is that, with $\hat{w}$ out of the way, the points $\overleftarrow{w}$ and $\vec{w}$ can be separated by open sets. The consequence is just what you'd expect.

**Proposition 3.3.E.** *The fractured interval $\overleftrightarrow{I}$ is Hausdorff.*

*Proof.* Pick two points $s < t$ in $\overleftrightarrow{I}$. If $\pi s \neq \pi t$, we can find disjoint neighborhoods of $\pi s$ and $\pi t$ in $I$ and pull them back to $\overleftrightarrow{I}$. If $\pi s = \pi t$, then $s = \overleftarrow{w}$ and $t = \vec{w}$ for some $w \in W$. Hence, $(-\infty, \overleftarrow{w}]$ and $[\vec{w}, \infty)$ are disjoint neighborhoods of $s$ and $t$. □

When studying $\overleftrightarrow{I}$, we found it useful to work with basis intervals whose leftmost and rightmost elements had been removed. For studying $\overleftrightarrow{I}$, it will be useful to go the opposite direction. Let's say a basis interval is *full* if it has both a leftmost element and a rightmost element. As we saw earlier, the full intervals in $\overleftrightarrow{I}$ are the ones that look like $[\vec{a}, \overleftarrow{b}]$. The full intervals in $\overleftrightarrow{I}$ are the same.

**Proposition 3.3.F.** *In $\overleftrightarrow{I}$, every full interval is compact.*

*Proof.* Consider a full interval $[\vec{a}, \overleftarrow{b}]$. Let $W'$ be the subset of $W$ lying between $\vec{a}$ and $\overleftarrow{b}$. Pick a function $\kappa\colon W' \to \mathbb{R}_+$ for which the sum $K = \sum_{w \in W'} \kappa w$ is finite, and let $\theta$ be the map from $[\vec{a}, \overleftarrow{b}] \cap \overleftrightarrow{I}$ to $\mathbb{R}$ given by the formula

$$\theta s = \pi s + \sum_{\substack{w \in W' \\ \hat{w} < s}} \kappa w.$$

It's not hard to see that $\theta$ is a homeomorphism whose image is the set

$$[a, b + K] \smallsetminus \bigcup_{w \in W'} (\theta\overleftarrow{w}, \theta\vec{w}),$$

which is closed and bounded. □

**Proposition 3.3.G.** *In $\overleftrightarrow{I}$, every full interval is clopen.*

*Proof.* In $\overleftrightarrow{I}$, the complement of a full interval $[\vec{a}, \overleftarrow{b}]$ is the union of the basis intervals $(-\infty, \overleftarrow{a}]$ and $[\vec{b}, \infty)$. □

**Proposition 3.3.H.** *If $W$ is dense in $I$, the full intervals form a basis for $\overleftrightarrow{I}$.*

*Proof.* Suppose $W$ is dense in $I$. Pick any point $s \in \overleftrightarrow{I}$ and any basis interval $(a, b)$ containing it. If $\pi a \neq \pi s$, find a point of $W$ in the interval $(\pi a, \pi s)$ and call it $\alpha$. If $\pi a = \pi s$, observe that $a = \hat{\alpha}$ and $s = \vec{\alpha}$ for some $\alpha \in W$. One way or another, we've found a point $\alpha \in W$ with $a < \vec{\alpha} \leq s$. Using the same technique, we can find a point $\beta \in W$ with $s \leq \overleftarrow{\beta} < b$. The full interval $[\vec{\alpha}, \overleftarrow{\beta}]$ is a neighborhood of $s$ contained in $(a, b)$. □

**Corollary 3.3.I.** *If $W$ is dense in $I$, every full interval in $\overleftrightarrow{I}$ is a Cantor set.*

*Proof.* Suppose $W$ is dense in $I$. Because we require $W$ to be countable or smaller, the results above imply that $\overleftrightarrow{I}$ is a Hausdorff space with a countable basis of clopen sets. Any full interval $H \subset \overleftrightarrow{I}$ has the same properties, and in addition is compact. Therefore, $H$ is a Cantor set as long as it has no isolated points [14, Theorem 3].

Intersecting $H$ with a basis interval in $\overleftrightarrow{I}$ yields another basis interval. Since every basis interval contains more than one point, it follows that $H$ has no isolated points. □

### 3.3.5 Metrization

The topology of $I$ is induced by the metric that $I$ inherits from $\mathbb{R}$. If $W$ is dense in $I$, the topology of $\overleftrightarrow{I}$ can be metrized too, and there's a simple way to do it. For the examples given in Section 3.3.2, the resulting metric is dynamically meaningful.

For the rest of this section, suppose $W$ is dense in $I$. Let's say we've assigned each point in $W$ a natural number, its *grade*, and there are only finitely many points of each grade. Since $W$ is countable or smaller, this is always possible.

For the binary shift, the points in $W$ are the numbers whose binary expansions are eventually constant, and they're naturally graded by the number of digits before the constant tail. For an interval exchange, the points in $W$ are the points that will eventually fall into a break under forward or backward iteration, and they're naturally graded by how long it takes for that to happen. For the sake of concreteness in our proof of Proposition 3.5.H, we'll fix a normalization by declaring the break points to have grade zero.

Pick a real number $K > 1$, and define the *height* of a point in $w \in W$ to be $K^{-\operatorname{grade} w}$. In $\overleftrightarrow{I}$, define the height of $\hat{w} \in \iota W$ to be the height of $w$, and the height of any other point to be zero. The heights of some points in the divided interval for the binary shift are illustrated below.

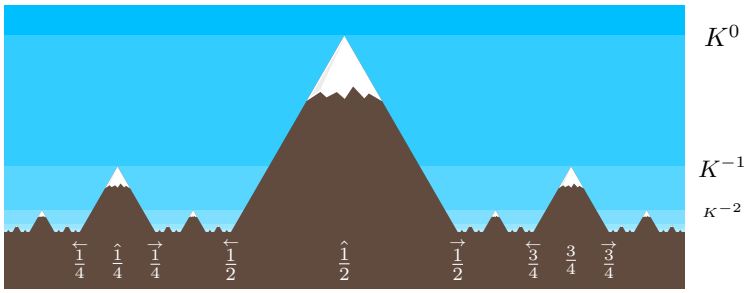


Let's say the distance between two points $a, b \in \overleftrightarrow{I}$ is the height of the highest point in $(a, b) \subset \overleftrightarrow{I}$. This defines a metric (in fact, an ultrametric) on $\overleftrightarrow{I}$, which I'll call the *division metric* with steepness $K$. The assumption that $W$ is dense in $I$ is essential here: it guarantees that distances between distinct points are positive.

**Proposition 3.3.J.** *The division metric induces the topology of* $\overleftrightarrow{I}$.

*Proof.* Let's see what the open balls of the division metric look like. Given a radius $r > 0$, let $W_{\geq r}$ be the set of points in $W$ with heights greater than or equal to $r$. This set is finite, because there are only finitely many points of each grade. Listing the points in $W_{\geq r}$ from left to right as $w_1, \ldots, w_n$, we can write down all the open balls of radius $r$:

$$(-\infty, \hat{w}_1), (\hat{w}_1, \hat{w}_2), \ldots, (\hat{w}_{n-1}, \hat{w}_n), (\hat{w}_n, \infty).$$

From this description, it's clear that the open balls of the division metric are open subsets of $\overset{\leftrightarrow}{I}$. It's also clear that every full interval is a union of open balls, because every full interval can be written as $(\hat{a}, \hat{b})$ for $a, b \in W$. By Proposition 3.3.H, the full intervals form a basis for $\overset{\leftrightarrow}{I}$, so we're done. □

For both of the examples in Section 3.3.2, the division metric tells you how long two points in the fractured interval travel together before they end up on opposite sides of a break. Nearby points move together for a long time, while the most distant points are separated immediately. A single step of the dynamics can take a pair of points at most one step closer to being separated, increasing the distance between then by at most a factor of $K$. That means the dynamical map is Lipschitz with respect to the division metric, essentially by construction.

The division metric on $\overset{\leftrightarrow}{I}$ and the Euclidean metric on $I$ are very different, so playing them off against one another might lead to amusing results. Let's get them talking by defining $\text{gap}_r$ to be the minimum distance between points in $W_{\geq r}$, according to the Euclidean metric on $I$. We can force the two metrics to work together by putting conditions on the gap function. Now, what sort of trouble can we start?

The division metric entertains all kinds of Hölder functions, but the Euclidean metric will not allow any interesting function to have a Hölder exponent greater than one. A function $f$ on $\overset{\leftrightarrow}{I}$ factors through $\pi$ if and only if its values match at adjacent edge points, in the sense that $f\overleftarrow{w} = f\vec{w}$ for all $w \in W$. If a Hölder function on $\overset{\leftrightarrow}{I}$ were to factor through $\pi$, there might be a bit of a problem.

**Theorem 3.3.K.** *Suppose the gap function falls off slower than a power law, so* $\text{gap}_r \geq Mr^\alpha$ *for some positive constants* $M$ *and* $\alpha$*. Consider a function* $f$ *from* $\overset{\leftrightarrow}{I}$ *into some metric space* $X$ *which factors through* $\pi$ *as shown:*

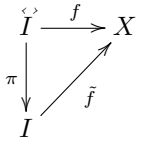


*If* $f$ *is Hölder with exponent* $\nu > 0$*, then* $\tilde{f}$ *is Hölder with exponent* $\nu/\alpha$*.*

Interesting, but not terribly entertaining. Let's turn up the heat.

**Corollary 3.3.L.** *Suppose the gap function falls off slower than* every *power law: for any* $\alpha > 0$*, no matter how small, we can find a positive constant* $M$ *with*

$\text{gap}_r \geq Mr^\alpha$. *If $f$ is a Hölder function on $\hat{I}$ whose values match at adjacent endpoints, then $f$ is constant.*

A bit of terminology will speed our proof of Theorem 3.3.K. Let's say a point $s \in \hat{I}$ is in the *left watershed* of $w \in W$ if $s < \hat{w}$, and every point in $(s, \hat{w}) \subset \overleftrightarrow{I}$ is lower than $w$. Define the *right watershed* similarly. The left and right watersheds of $\frac{1}{4}$ and $\frac{5}{8}$ for the binary shift are sketched below.

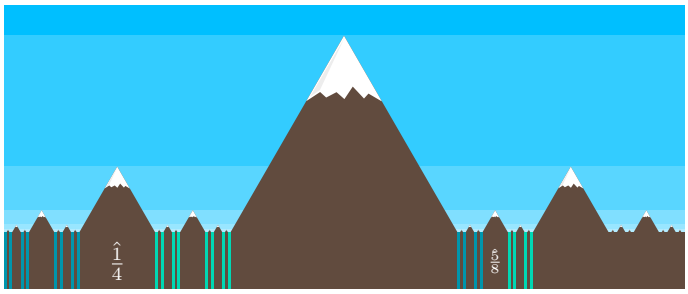


Watersheds are useful because of the following fact.

**Lemma 3.3.M.** *For any $s \in I$, if $\vec{\iota}s$ is in the left watershed of $w \in W$, then $\text{gap}_{d(\vec{\iota}s, \overleftarrow{w})} \leq w - s$. The mirror image statement also holds.*

*Proof.* Let $\hat{v} \in \iota W$ be the highest point in $(\vec{\iota}s, \hat{w}) \subset \overleftrightarrow{I}$, and let $h$ be its height. Since $\vec{\iota}s$ is in the left watershed of $w$, the point $w$ is higher than $v$. That means $v$ and $w$ are both in $W_{\geq h}$, so $w - v$ is at least $\text{gap}_h$. Thus, $w - s$ is at least $\text{gap}_h$. Since $\hat{v}$ was the highest point between $\vec{\iota}s$ and $\hat{w}$, we have $d(\vec{\iota}s, \overleftarrow{w}) = h$ by definition, so we're done. The mirror image statement is proven similarly. □

*Proof of Theorem 3.3.K.* Suppose $f$ is Hölder with exponent $\nu > 0$ and scale constant $C > 0$. To see that $\tilde{f}$ is Hölder, pick any two points $a < b$ in $I$. Observe that

$$d(\tilde{f}a, \tilde{f}b) = d(f\vec{\iota}a, f\overleftarrow{\iota}b).$$

Let $w \in W$ be the highest point in $(\vec{\iota}a, \overleftarrow{\iota}b) \subset \overleftrightarrow{I}$. By assumption, $f\overleftarrow{w} = f\vec{w}$, so the triangle inequality and the Hölder condition on $f$ give

$$\begin{aligned} d(\tilde{f}a, \tilde{f}b) &\leq d(f\vec{\iota}a, f\overleftarrow{w}) + d(f\vec{w}, f\overleftarrow{\iota}b) \\ &\leq Cd(\vec{\iota}a, \overleftarrow{w})^\nu + Cd(\vec{w}, \overleftarrow{\iota}b)^\nu. \end{aligned}$$

Since $\vec{\iota}a$ is in the left watershed of $w$, the lemma tells us that

$$\text{gap}_{d(\vec{\iota}a, \overleftarrow{w})} \leq w - a \qquad\qquad \text{gap}_{d(\vec{w}, \overleftarrow{\iota}b)} \leq b - w.$$

Our lower bound on the gap function then ensures that

$$Md(\vec{\iota}a, \overleftarrow{w})^\alpha \leq w - a \qquad\qquad Md(\vec{w}, \overleftarrow{\iota}b)^\alpha \leq b - w,$$

so

$$d(\tilde{f}a, \tilde{f}b) \leq B[(w-a)^{\nu/\alpha} + (b-w)^{\nu/\alpha}]$$

for $B = \frac{C}{M^{\nu/\alpha}}$. Since the function $t \mapsto t^{\nu/\alpha}$ is non-decreasing, it follows[4] that

$$\begin{aligned} d(\tilde{f}a, \tilde{f}b) &\leq 2B[(w-a) + (b-w)]^{\nu/\alpha} \\ &= 2B(b-a)^{\nu/\alpha}. \end{aligned}$$

Since $a$ and $b$ were arbitrary, aside from the condition that $a < b$, we've shown that $\tilde{f}$ is Hölder with exponent $\nu/\alpha$ and scale constant $\frac{2C}{M^{\nu/\alpha}}$. □

## 3.4 Dividing translation surfaces

### 3.4.1 Construction of divided and fractured surfaces

Now, let's move up to the two-dimensional case. Cover $\Sigma \smallsetminus \mathfrak{B}$ with flow boxes. Each flow box can be identified with a rectangle $I \times L \subset \mathbb{R}^2$, where $I$ and $L$ are open intervals in $\mathbb{R}$. The critical leaves of $\Sigma$ intersect the flow box as vertical lines $\{w\} \times L$. There are only finitely many critical leaves, and each one passes through the flow box at most countably many times. Dividing $I$ at the positions of the critical leaves, we can produce a *divided flow box* $\overleftrightarrow{I} \times L$. In the divided flow box, each critical leaf $\{w\} \times L$ splits into a *left lane* $\{\overleftarrow{w}\} \times L$, a *median* $\{\hat{w}\} \times L$, and a *right lane* $\{\vec{w}\} \times L$.

The transition maps between the flow boxes induce transitions between the divided flow boxes in a natural way. Gluing the divided flow boxes together along these transitions yields a new space—a divided version of $\Sigma \smallsetminus \mathfrak{B}$.

We still need to decide what to do with the singularities. The region around a singularity $\mathfrak{b} \in \mathfrak{B}$ can be built from cut square pieces, as described in Section 3.2.1. Here's one of them, with the critical leaves marked. (Since the critical leaves are generally dense in $\Sigma$, I've only drawn finitely many of them thick enough to see.)

The leaf through the center point is critical, of course, because the center point is $\mathfrak{b}$ itself. Here's a divided version of the same cut square:

[4] If $\rho\colon [0,\infty) \to \mathbb{R}$ is non-decreasing, and $t \leq t'$,

$$\rho t + \rho t' \leq 2\rho t' \leq 2\rho(t+t').$$

Many thanks to Sona Akopian for pointing this out.

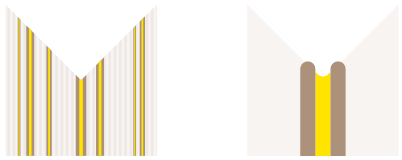

Zooming in, you can see that I've cut away $\hat{\mathfrak{b}}$, but kept $\overleftarrow{\mathfrak{b}}$ and $\overrightarrow{\mathfrak{b}}$. When a singularity is built from pieces that look like this one (and its half-rotation), the left and right lanes of the center leaf will connect up into lines running past the singularity, but the medians will remain disconnected rays that end at the singularity. Topologically, the region around a divided singularity looks like this:

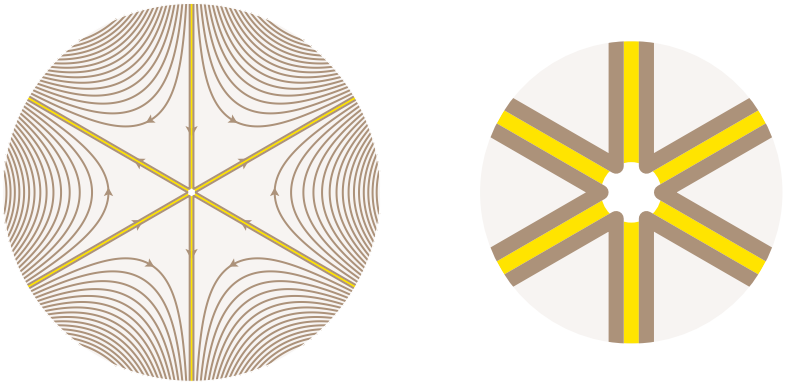

The zoomed-in picture on the right shows how the lanes of the critical leaves join up around the singularity. For clarity, only the parts of the critical leaves adjacent to the singularity are shown.

Now that we've decided what to do with the singularities, we can extend our divided version of $\Sigma \smallsetminus \mathfrak{B}$ to a full *divided surface* $\overleftrightarrow{\Sigma}$. The quotient map discussed in the one-dimensional case extrudes naturally to a quotient map $\pi\colon \overleftrightarrow{\Sigma} \to \Sigma$. Away from the singularities, the embedding from the one-dimensional case also extrudes, yielding an embedding $\Sigma \smallsetminus \mathfrak{B} \to \overleftrightarrow{\Sigma}$. This embedding can't be extended over the singularities, because a singularity has no median point associated with it, and sending it to an associated lane point would break continuity.

We'll refer to the $\pi$-preimage of a leaf of $\Sigma$ as a *road*.[5] If $\mathcal{L}$ is a critical leaf of $\Sigma$, the *critical road* $\pi^{-1}\mathcal{L}$ splits into a left lane $\{\overleftarrow{w} : w \in \mathcal{L}\}$, a median $\{\hat{w} : w \in \mathcal{L}\}$, and a right lane $\{\overrightarrow{w} : w \in \mathcal{L}\}$, as we saw locally at the beginning of the section. Each of the lane points associated with a singularity is adjacent to a lane of a forward-critical road and a lane of a backward-critical road. For convenience, we'll consider it an honorary member of both.

[5] Our one-way roads are technically different from, but morally related to, the *two-way streets* of [5].

As in the one-dimensional case, define the *fractured surface* $\overset{\leftrightarrow}{\Sigma}$ to be the complement of the critical leaves of $\Sigma$—or, more precisely, their images under $\iota$—in the divided surface $\overset{\leftrightarrow}{\Sigma}$. The divided flow boxes and divided singularity charts from which we constructed $\overset{\leftrightarrow}{\Sigma}$ pull back to an atlas of *fractured flow boxes* and *fractured singularity charts* on $\overset{\leftrightarrow}{\Sigma}$.

### 3.4.2 Properties of divided and fractured surfaces

Within each divided flow box $\overset{\leftrightarrow}{I} \times L$, we can find more flow boxes of the form $H \times L$, where $H \subset \overset{\leftrightarrow}{I}$ is a basis interval. For convenience, let's define the notation $\overset{\leftrightarrow}{H} = H \cap \overset{\leftrightarrow}{I}$. A flow box $U = H \times L$ will be called *full* if $H$ is full, and *well-cut* if $H$ is full and $\iota^{-1}U$ is well-cut. Define $\overset{\leftrightarrow}{U}$ as $U \cap \overset{\leftrightarrow}{\Sigma}$, or equivalently $\overset{\leftrightarrow}{H} \times L$, and $\operatorname{trim} U$ as $(\operatorname{trim} H) \times L$. It's apparent from the analogous one-dimensional result that $\pi^{-1}\iota^{-1}U = \operatorname{trim} U$ for any flow box $U$.

Flow boxes form a basis for $\overset{\leftrightarrow}{\Sigma}$. If the critical leaves are dense in $\Sigma$, full flow boxes form a basis for $\overset{\leftrightarrow}{\Sigma}$, as a consequence of Proposition 3.3.H. If full flow boxes form a basis for $\overset{\leftrightarrow}{\Sigma}$, well-cut flow boxes do too, as a consequence of Proposition 3.2.B.

Propositions 3.3.A, 3.3.B, and 3.3.E carry over from dimension one straightforwardly enough that I'll state them without proof.

**Proposition 3.4.A.** *The embedding of $\Sigma$ in $\overset{\leftrightarrow}{\Sigma}$ is dense.*

**Proposition 3.4.B.** *The divided surface $\overset{\leftrightarrow}{\Sigma}$ is locally connected.*

**Proposition 3.4.C.** *The fractured surface $\overset{\leftrightarrow}{\Sigma}$ is Hausdorff.*

Because $\Sigma$ is compact, Proposition 3.3.F can be extended to a global result.

**Proposition 3.4.D.** *The fractured surface $\overset{\leftrightarrow}{\Sigma}$ is compact.*

*Proof.* Because $\Sigma$ compact, it can be covered by a finite collection of flow boxes and singularity charts. Therefore, $\overset{\leftrightarrow}{\Sigma}$ can be covered by a finite collection of fractured flow boxes and fractured singularity charts. Each fractured flow box $\overset{\leftrightarrow}{I} \times L$ can be covered almost out to the edges by a "full closed box" of the form $\overset{\leftrightarrow}{H} \times C$, where $H \subset \overset{\leftrightarrow}{I}$ is full interval and $C \subset L$ is a closed interval. Similarly, each fractured singularity chart can be covered almost out to the edges by a finite collection of full closed boxes, taking crucial advantage of the fact that the non-compact medians of the critical rays have been removed.

With a little care, we can now cover $\overset{\leftrightarrow}{\Sigma}$ with a finite collection of full closed boxes. Each full closed box is a product of two compact spaces, so we've covered $\overset{\leftrightarrow}{\Sigma}$ with a finite collection of compact sets. □

Theorem 3.3.C carries over with a catch: dividing a translation surface opens up a tiny hole at each singularity, and a local system on the divided surface can have nontrivial holonomies around those holes.

**Theorem 3.4.E.** *For any linear group $G$, the direct image functors $\pi_*$ and $\iota_*$ give an equivalence between the category of $G$ local systems on $\overleftrightarrow{\Sigma}$ and the category of $G$ local systems on $\Sigma \smallsetminus \mathfrak{B}$.*

The reason for the theorem remains the same.

**Lemma 3.4.F.** *If $\mathcal{F}$ is a local system on $\overleftrightarrow{\Sigma}$, the restriction $\mathcal{F}_{\operatorname{trim} U \subset U}$ is an isomorphism for any flow box $U$.*

*Proof.* Write $U$ as $H \times L$ for some basis interval $H \subset \overleftrightarrow{I}$ and repeat the proof of Lemma 3.3.D, replacing $A$ with $A \times L$. □

*Proof of Theorem 3.4.E.* Replace

$$\begin{array}{ccc} & \overleftrightarrow{I} & \\ \pi \swarrow & & \nwarrow \iota \\ I & & \end{array} \quad \text{with} \quad \begin{array}{ccc} & \overleftrightarrow{\Sigma} & \\ \pi \swarrow & & \nwarrow \iota \\ \Sigma \smallsetminus \mathfrak{B} & & \end{array}$$

in the proof of Theorem 3.3.C, and change all the basis intervals to flow boxes. □

## 3.5 Dynamics on divided surfaces

### 3.5.1 The vertical flow

Within a flow box $I \times L$ in $\Sigma$, the vertical flow is easy to describe:

$$\psi^t \colon (s, \zeta) \mapsto (s, \zeta + t),$$

as long as $t$ is small enough that $\zeta + t$ is still in the interval $L$. This local description can be lifted directly to the divided flow box $\overleftrightarrow{I} \times L$ in $\overleftrightarrow{\Sigma}$; the only change is that $s$ will now be a point in $\overleftrightarrow{I}$ instead of in $I$.

The vertical flow on a singularity chart can be lifted in the same way. It's worth thinking carefully about how the lifted flow acts on the critical roads of $\overleftrightarrow{\Sigma}$. Pick a singularity chart on $\Sigma$, and look at a point $w$ on a forward-critical leaf that plunges into the singularity without leaving the chart. In $\overleftrightarrow{\Sigma}$, the point $w$ splits into the three points $\overleftarrow{w}, \hat{w}, \vec{w}$ on a forward-critical road. As time runs forward, the point $w$ falls into the singularity and disappears. Its lift $\hat{w}$, on the median of the road, does the same. The points $\overleftarrow{w}$ and $\vec{w}$, however, follow the left and right lanes past the singularity, peeling off in different directions:

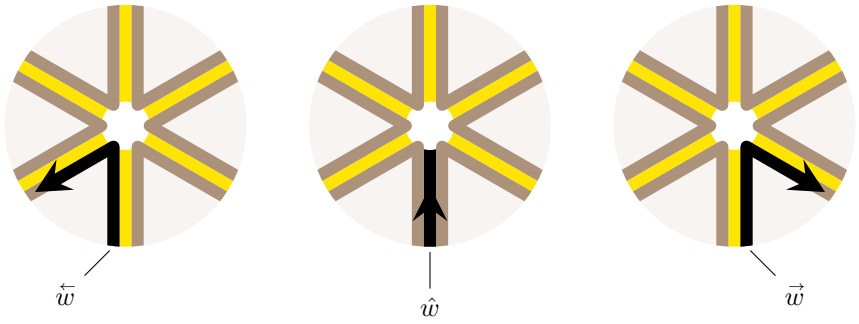


On a backward-critical road, the story is the same, but told backward.

Just as the local vertical flows on flow boxes and singularity charts fit together into a global vertical flow on $\Sigma \smallsetminus \mathfrak{B}$, their lifts fit together into a global vertical flow on $\overleftrightarrow{\Sigma}$. For convenience, I'll refer to this vertical flow also as $\psi$. The vertical flows on $\overleftrightarrow{\Sigma}$ and $\Sigma \smallsetminus \mathfrak{B}$ commute with the embedding $\iota$. They don't quite commute with the quotient map $\pi$, because when a point on a critical leaf of $\Sigma \smallsetminus \mathfrak{B}$ falls into a singularity, its left- and right-lane lifts keep going. The forward vertical flows on $\Sigma \smallsetminus \mathfrak{W}^-$ and its $\pi$-preimage, however, do commute with $\pi$. The same goes for the backward vertical flows on $\Sigma \smallsetminus \mathfrak{W}^+$ and its $\pi$-preimage.

The vertical flow on $\overleftrightarrow{\Sigma}$, like the one on $\Sigma \smallsetminus \mathfrak{B}$, is a flow by bicontinuous relations. This should not be taken for granted: the topology of a divided interval was engineered to make it so. The medians of the critical leaves are the only parts of $\overleftrightarrow{\Sigma}$ that vanish into the singularities under the vertical flow. Removing them leaves an invariant subspace, $\overleftrightarrow{\Sigma}$, on which the vertical flow acts by homeomorphisms. Thus, while the vertical flows on $\Sigma \smallsetminus \mathfrak{B}$ and $\overleftrightarrow{\Sigma}$ may look a bit ugly, the vertical flow on $\overleftrightarrow{\Sigma}$ is remarkably well-behaved: it's a flow by homeomorphisms on a compact Hausdorff space.

The divided and fractured surfaces are foliated by the orbits of the vertical flow—or they would be, at least, if foliations were usually defined on more general spaces than manifolds. With that in mind, we'll sometimes refer to the orbits of the vertical flows on $\overleftrightarrow{\Sigma}$ and $\overleftrightarrow{\Sigma}$ as vertical leaves.

### 3.5.2 First return maps

Let's say a *horizontal segment* on a divided surface is a subset that looks like a horizontal basis interval in some flow box. More precisely, it's a horizontal slice $H \times \{\zeta\}$ across a divided flow box $H \times L$. The quotient map $\pi$ projects horizontal segments on $\overleftrightarrow{\Sigma}$ down to horizontal segments on $\Sigma$, as defined in Section 3.2.2. Lemma 3.2.A, which made it sensible to talk about first return maps on a translation surface, has an analogue on the divided surface.

**Lemma 3.5.A.** *Let $Z$ be a horizontal segment on $\overleftrightarrow{\Sigma}$, and let $p$ be a point in $Z$. Unless $p$ is on the median of a forward-critical road, the vertical flow will*

*eventually carry $p$ back to $Z$.*

The proof is edifying, but rather long, so I've postponed it to the end of the section. The most important consequence of this lemma is that we can define a first return relation on any horizontal segment in $\overleftrightarrow{\Sigma}$, just like we did for horizontal segments in $\Sigma$.

Suppose $Z$ is a horizontal slice across a well-cut flow box $U = H \times L$ in $\overleftrightarrow{\Sigma}$. Then $\iota^{-1}Z$ is a well-cut horizontal segment in $\Sigma$, so the first return relation on $\iota^{-1}Z$ is an interval exchange. The divided version of that interval exchange, constructed as in Section 3.3.2, is precisely the first return relation on $Z$.

Identifying $Z$ and $\iota^{-1}Z$ with $H$ and $\iota^{-1}H$ in the obvious way, we can think of the first return relations as relations on $H$ and $\iota^{-1}H$. These relations don't depend on which slice across $H \times L$ we take, so we can talk about the first return relations on $H$ and $\iota^{-1}H$ defined by a well-cut flow box $H \times L$ without picking a slice at all. The first return relation becomes a function when restricted to $\dot{\dot{H}}$.

We'll be working with these first return relations a lot, so it will be useful to set down a pattern of notation for talking about them. For convenience, I'll refer to the first return relations on $H$ and $\iota^{-1}H$ both as $\alpha$. Let $W \subset \iota^{-1}H$ be the positions of the critical leaves, recalling that $\dot{\dot{H}} = H \smallsetminus \iota W$. Within $W$, let $W^+$ and $W^-$ be the positions of the backward- and forward-critical leaves. If $\Sigma$ has no saddle connections, $W^+$ and $W^-$ are disjoint, and thus form a partition of $W$.

Let $B^+ \subset W^+$ be the places where the backward-critical leaves first pass through $\iota^{-1}U$ after shooting out of their singularities. Similarly, let $B^- \subset W^-$ be the places where the forward-critical leaves last pass through $\iota^{-1}U$ before diving into their singularities. The sets $B^+$ and $B^-$ are the break points of the backward and forward first return relations $\alpha^{-1}$ and $\alpha$ on $\iota^{-1}H$, as described in Section 3.3.2.

*Proof of Lemma 3.5.A.* If $p$ is in $\iota\Sigma$, we can just apply Lemma 3.2.A to an appropriate closed subinterval of $\pi Z$, and we're done. The only time we have to do something less direct is when $p$ is in the left or right lane of a critical road.

Suppose $p = \vec{w}$ for some $w \in \mathfrak{W} \smallsetminus \mathfrak{W}^-$. Because it's in the right lane of a critical road, $p$ can't be the rightmost element of a horizontal basis interval. We can therefore assume, without loss of generality, that $Z$ has no rightmost element. In this case, the only way for $\pi^{-1}\pi Z$ to contain more points than $Z$ is for $Z$ to have a leftmost point $\vec{a}$, in which case $\pi^{-1}\pi Z = \{\overleftarrow{a}, \hat{a}\} \cup Z$.

By Lemma 3.2.A, $\psi^t w$ returns to $\pi Z$ at some time $t > 0$. Away from $\mathfrak{W}^-$ and its $\pi$-preimage, the forward vertical flows on $\overleftrightarrow{\Sigma}$ and $\Sigma$ commute, so $\pi\psi^t p \in \pi Z$. In other words, $\psi^t p \in \pi^{-1}\pi Z$. Points traveling along the vertical flow on $\overleftrightarrow{\Sigma}$ never change lanes, so $\psi^t p$ is in a right lane. Hence, knowing that $\psi^t p \in \pi^{-1}\pi Z$, we can conclude that $\psi^t p \in Z$, which is what we wanted to show. The same reasoning can be used when $p = \overleftarrow{w}$ for some $w \in \mathfrak{W} \smallsetminus \mathfrak{W}^-$.

On the other hand, suppose $p = \vec{w}$ for some $w \in \mathfrak{W}^-$. In this case, we can assume without loss of generality that $p$ is the leftmost point of $Z$. Follow $\psi^t Z$

upward along the vertical flow until $\psi^t p$ passes a singularity, exiting the forward-critical road it started on and merging onto the adjacent backward-critical road. By this time, a few pieces of $\psi^t Z$ may have already hit singularities and peeled off to the right, but some piece of $\psi^t Z$ is still traveling with $\psi^t p$.

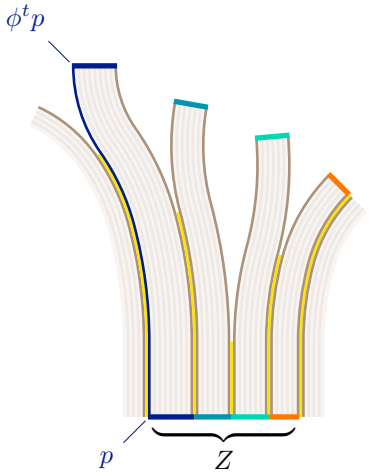


The backward-critical road $\psi^t p$ is now following might also be forward-critical—a saddle connection. In that case, we can repeat the same maneuver. Because $\Sigma$ is compact, it has only finitely many forward-critical roads, so $\psi^t p$ will eventually end up on either a road that isn't forward-critical or a road that it's traveled before. In the latter case, because $\psi^t p$ can only merge onto a road at its very beginning, $\psi^t p$ will eventually return to a point it's passed through before. Since $\psi$ is a flow by bicontinuous relations, as defined in Appendix B, it follows that $\psi^t p$ is defined for all $t \in \mathbb{R}$, and periodic. As a result, $\psi^t p$ will eventually return to its starting point, and thus to $Z$.

This leaves us with the case where $\psi^t p$ finally merges onto a road that isn't forward critical. In other words, at some time $t > 0$, we have $\pi\psi^t p \in \mathfrak{W} \smallsetminus \mathfrak{W}^-$. Let $Y$ be the piece of $\psi^t Z$ that's stuck with $\psi^t p$ all this time. We showed earlier that the vertical flow eventually carries $\psi^t p$ back to $Y$. The vertical flow is injective, so $\psi^t p$ must have passed through $\psi^{-t} Y$ on its way back to $Y$. Since $\psi^{-t} Y \subset Z$, we've proven that $\psi^t p$ eventually returns to $Z$. □

### 3.5.3 Minimality

Assuming that the vertical flow on $\overset{\circ}{\Sigma}$ is minimal will make our discussion of abelianization and the conditions required for it go much more smoothly—so much so that I suspect it's a necessary condition for abelianization to make sense. Fortunately, this condition follows from a simple condition on $\Sigma$.

**Proposition 3.5.B.** *If $\Sigma$ has no saddle connections, the forward and backward vertical flows on $\overset{\circ}{\Sigma}$ are both minimal as dynamical systems.*

*Proof.* The forward and backward cases are mirror images of each other, so let's focus on the forward vertical flow. Suppose $\Sigma$ has no saddle connections, recalling that this implies $\Sigma$ is minimal. We want to prove that the forward orbit $\mathcal{P}$ of any point $p \in \overleftrightarrow{\Sigma}$ is dense.

Suppose $p = \iota q$ for some $q \in \Sigma \smallsetminus \mathfrak{W}$. Then we can just observe that the minimality of $\Sigma$ implies that the forward orbit of $q$ is dense in $\Sigma$. Since $\iota\Sigma$ is dense in $\overleftrightarrow{\Sigma}$ by Proposition 3.4.A, we're done.

Suppose $p$ is in the left or right lane of a backward-critical road. In other words, $p \in \{\overleftarrow{w}, \overrightarrow{w}\}$ for some $w \in \mathfrak{W}^+$. Let $\mathcal{W}$ be the forward orbit of $w$. Because $\Sigma$ has no saddle connections, the backward-critical leaf $w$ lies on can't also be forward-critical, so the minimality of $\Sigma$ implies that $\mathcal{W}$ is dense in $\Sigma$.

Because $\Sigma$ is minimal, full flow boxes form a basis for $\overleftrightarrow{\Sigma}$, as pointed out in Section 3.4.2. Thus, to show that $\mathcal{P}$ is dense in $\overleftrightarrow{\Sigma}$, we just have to show that it intersects every full flow box $\overleftrightarrow{U}$. Note that $U$ is a full flow box in $\overleftrightarrow{\Sigma}$. Since $\iota$ is continuous, $\iota^{-1}U$ is an open subset of $\Sigma$, and therefore intersects the dense orbit $\mathcal{W}$. That means $\pi^{-1}\iota^{-1}U$ intersects $\iota\mathcal{W}$. Recall from Section 3.4.2 that $\pi^{-1}\iota^{-1}U = \operatorname{trim} U$. In $\overleftrightarrow{\Sigma}$, any open subset that contains a point in the median of a critical road also contains the corresponding points in the left and right lanes, so $\operatorname{trim} U$ intersects $\mathcal{P}$. Since $\operatorname{trim} U \subset U$, we've shown that $U$ intersects $\mathcal{P}$. Since $\mathcal{P} \subset \overleftrightarrow{\Sigma}$, it follows that $\overleftrightarrow{U}$ intersects $\mathcal{P}$. Since $\overleftrightarrow{U}$ could have been any full flow box in $\overleftrightarrow{\Sigma}$, we've proven that $\mathcal{P}$ is dense in $\overleftrightarrow{\Sigma}$, under the assumption that $p$ is in the left or right lane of a backward-critical road.

Finally, suppose $p$ is in the left or right lane of a forward-critical road. The forward vertical flow will eventually carry $p$ past a singularity, where it will leave its current road and merge onto an adjacent backward-critical road. That means the forward orbit of $p$ contains the left or right lane of a backward-critical road. We just proved that the left and right lanes of a backward-critical road are both dense in $\overleftrightarrow{\Sigma}$, so we're done. □

On the fractured surface, a minimal vertical flow induces minimal first return maps.

**Proposition 3.5.C.** *Let $H \times L$ be a well-cut flow box in $\overleftrightarrow{\Sigma}$. If the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal, the first return map on $\overleftrightarrow{H}$ is too.*

*Proof.* Suppose the first return map on $\overleftrightarrow{H}$ is not minimal. Find a closed invariant subset $C \subset \overleftrightarrow{H}$ which is neither empty nor all of $\overleftrightarrow{H}$. Let $\mathcal{C}$ be the orbit of $C \times L$ under the vertical flow. Notice that $\mathcal{C}$ can't intersect the open set $(\overleftrightarrow{H} \smallsetminus C) \times L$: if it did, the first return map on $\overleftrightarrow{H}$ would send some element of $C$ into $\overleftrightarrow{H} \smallsetminus C$. Since $\mathcal{C}$ is made of vertical orbits, we've shown that not every orbit of the vertical flow is dense in $\overleftrightarrow{\Sigma}$. Hence, the vertical flow on $\overleftrightarrow{\Sigma}$ is not minimal. □

### 3.5.4 Ergodic theory

Proposition 3.4.A says that $\overleftrightarrow{\Sigma}$ is no bigger than $\Sigma$ with respect to continuous functions, and Theorem 3.4.E says the same with respect to locally constant

sheaves. It will be useful to have a similar result with respect to vertically invariant measures. Such a thing should be true, because passing from $\Sigma \smallsetminus \mathfrak{W}$ to $\overleftrightarrow{\Sigma}$ just means adding finitely many vertical leaves, which ought to have measure zero.

To turn these intuitions into something tangible, first equip $\Sigma$ and $\overleftrightarrow{\Sigma}$ with their Borel $\sigma$-algebras. You can show, with a little thought, that any vertical slice across a flow box of $\overleftrightarrow{\Sigma}$ is measurable. It follows that every vertical leaf of $\overleftrightarrow{\Sigma}$ is measurable, allowing us to formulate the following proposition.

**Proposition 3.5.D.** *If the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal, an invariant probability measure on $\overleftrightarrow{\Sigma}$ assigns measure zero to every vertical leaf.*

*Proof.* Assume the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal. Let $\mathcal{L}$ be a vertical leaf. Pick a segment $\mathcal{L}_0$ of $\mathcal{L}$ that looks like $\{s\} \times [0, \tau)$ in some divided flow box $\vec{H} \times (-\tau, \tau)$. Observe that $\mathcal{L}_0$ is measurable. Let $\mathcal{L}_n = \psi^{n\tau}\mathcal{L}_0$ for each integer $n$. The segments $\mathcal{L}_n$ can't overlap locally, and a minimal flow on $\overleftrightarrow{\Sigma}$ has no periodic orbits, so in fact all the segments are disjoint.

Let $\mu$ be an invariant probability measure on $\overleftrightarrow{\Sigma}$. The segments $\mathcal{L}_n$ tile $\mathcal{L}$, so

$$\mu\mathcal{L} = \sum_{n\in\mathbb{Z}} \mu\mathcal{L}_n.$$

Because it's invariant, $\mu$ assigns the same measure to each $\mathcal{L}_n$. Hence,

$$\mu\mathcal{L} = \sum_{n\in\mathbb{Z}} \mu\mathcal{L}_0.$$

For the sum on the right-hand side to converge, its value must be zero. □

**Corollary 3.5.E.** *If the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal, an invariant probability measure on $\overleftrightarrow{\Sigma}$ assigns measure one to $\iota(\Sigma \smallsetminus \mathfrak{W})$.*

*Proof.* The only parts of $\overleftrightarrow{\Sigma}$ that don't lie in $\iota(\Sigma \smallsetminus \mathfrak{W})$ are the vertical leaves containing the left and right lanes of the critical roads. There are only finitely many of these, and each one has measure zero by the previous proposition. □

The corollary above leads to the following proposition, which links the ergodic properties of $\overleftrightarrow{\Sigma}$ and $\Sigma \smallsetminus \mathfrak{W}$.

**Proposition 3.5.F.** *If the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal, then it's uniquely ergodic if and only if the vertical flow on $\Sigma \smallsetminus \mathfrak{W}$ is.*

*Proof.* Assume the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal. The pushforward of an ergodic measure along a morphism of dynamical systems is ergodic, and the restriction of an ergodic measure to a measure-one invariant subspace is also ergodic. Therefore, $\overleftrightarrow{\Sigma}$ has an ergodic measure if and only if $\Sigma \smallsetminus \mathfrak{W}$ does.

Suppose $\Sigma \smallsetminus \mathfrak{W}$ is not uniquely ergodic. Distinct ergodic measures on $\Sigma \smallsetminus \mathfrak{W}$ push forward along $\iota$ to distinct ergodic measures on $\overleftrightarrow{\Sigma}$, so $\overleftrightarrow{\Sigma}$ isn't uniquely ergodic either.

On the other hand, suppose $\Sigma \smallsetminus \mathfrak{W}$ is uniquely ergodic. Then every ergodic measure on $\overleftrightarrow{\Sigma}$ must restrict to the same measure on $\Sigma \smallsetminus \mathfrak{W}$. All ergodic measures on $\overleftrightarrow{\Sigma}$ assign $\iota(\Sigma \smallsetminus \mathfrak{W})$ measure one, so if they agree on $\iota(\Sigma \smallsetminus \mathfrak{W})$, they agree everywhere. Hence, $\overleftrightarrow{\Sigma}$ is uniquely ergodic. □

Just like minimality, unique ergodicity of the vertical flow on $\overleftrightarrow{\Sigma}$ can be localized to any well-cut flow box, as long as the vertical flow is minimal.

**Proposition 3.5.G.** *Let $H \times L$ be a well-cut flow box in $\overleftrightarrow{\Sigma}$. If the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal and uniquely ergodic, the first return map on $\vec{H}$ is too.*

The proof should bring to mind a fractured analogue of the idea that a minimal translation surface can be expressed as a suspension of the first return map on any well-cut segment [15]. We don't need the full power of this idea, however, so we won't develop it in detail.

*Proof.* Pick a horizontal slice $Z = H \times \{\zeta\}$ across $H \times L$. Removing the break points of the first return map from $Z$ turns it into a disjoint union of full basis intervals $Z_1, \ldots, Z_k$. As each segment $Z_j$ is carried back to $Z$ by the vertical flow, it sweeps out a "full half-open box" of the form $Y_j \cong Z_j \times [0, T_j)$. Because the vertical flow on $\overleftrightarrow{\Sigma}$ is minimal, $\vec{Y}_1, \ldots, \vec{Y}_k$ cover $\overleftrightarrow{\Sigma}$, forming a measurable partition of it.

A measure on $\vec{Z}$ restricts to a measure on each segment $\vec{Z}_j$. The measures on $\vec{Z}_1, \ldots, \vec{Z}_k$ induce product measures on $\vec{Y}_1, \ldots, \vec{Y}_k$, which fit together into a measure on $\overleftrightarrow{\Sigma}$. In this way, a measure $\mu$ on $\vec{Z}$ can be extruded to a measure $\tilde{\mu}$ on $\overleftrightarrow{\Sigma}$. If $\mu$ is invariant under the first return map, $\tilde{\mu}$ is invariant under the vertical flow. If $\mu$ is a probability measure, $\tilde{\mu}\overleftrightarrow{\Sigma}$ is finite, so we can normalize $\tilde{\mu}$ to get a probability measure. Different probability measures on $Z$ produce different probability measures on $\overleftrightarrow{\Sigma}$.

Let $\mu$ be an ergodic measure on $Z$, and $\tilde{\mu}$ the measure on $\overleftrightarrow{\Sigma}$ constructed from it. Consider an invariant subset $\mathcal{Q}$ of $\overleftrightarrow{\Sigma}$. Because $\mathcal{Q}$ is invariant, its intersection with the box $\vec{Y}_j$ looks like $Q_j \times [0, T_j)$, where $Q_j \subset Z_j$ is a measurable set. The $\tilde{\mu}$-measure of $\mathcal{Q}$ is given in terms of $\mu$ by

$$\tilde{\mu}\mathcal{Q} = T_1\, \mu Q_1 + \ldots + T_k\, \mu Q_k.$$

The disjoint union $Q = Q_1 \cup \ldots \cup Q_k$ is an invariant subset of $Z$, so its $\mu$-measure is either zero or one. It follows that $\tilde{\mu}\mathcal{Q}$ is either zero or $T_1\, \mu Z_1 + \ldots + T_k\, \mu Z_k$, which is equal to $\tilde{\mu}\overleftrightarrow{\Sigma}$. We now see that $\tilde{\mu}$ assigns each invariant subset either zero measure or full measure, so its normalization is an ergodic measure on $\overleftrightarrow{\Sigma}$.

We've shown that any ergodic measure on $Z$ can be used to produce an ergodic measure on $\overleftrightarrow{\Sigma}$, and we remarked earlier that different probability measures

on $Z$ produce different probability measures on $\overleftrightarrow{\Sigma}$. Hence, if the first return map on $Z$ isn't uniquely ergodic, the vertical flow on $\overleftrightarrow{\Sigma}$ can't be either. □

### 3.5.5 The fat gap condition

On a generic translation surface, you should expect the critical leaves to fill up each flow box more or less evenly as you follow them out from the singularities, rather than clumping together. The *fat gap condition* formalizes this property. When horizontal distances on the fractured surface are measured using the the division metric from Section 3.3.5, the fat gap condition implies the hypothesis of Corollary 3.3.L, which restricts the behavior of Hölder functions on horizontal segments. This control over horizontal Hölder functions will be the crux of our argument in Section 8 that abelianization does what it's supposed to.

For convenience, we'll keep track of first return relation break points in this section using the height function from Section 3.3.5, with steepness $K$. The definition we're about to present won't depend on the height function, but stating it in terms of the height function will allow for more consistent notation, and it will help us connect the fat gap condition with the properties of the division metric.

**The local definition** Consider a well-cut horizontal segment $Z$ in $\Sigma$. As usual, let $W \subset Z$ be the positions of the critical leaves. Recall from Section 3.3.5 that $W_{\geq r}$ is the set of points in $W$ with heights greater than or equal to $r$, and $\text{gap}_r$ is the minimum distance between points in $W_{\geq r}$ according to the Euclidean metric on $Z$. Let's say $Z$ satisfies the *fat gap condition* if for all $\lambda > 0$, no matter how small,

$$K^{\lambda n} \, \text{gap}_{K^{-n}}$$

is bounded below by a positive number as $n$ varies in $\mathbb{N}$. Intuitively, the fat gap condition says that the break points of the iterated first return relations $\alpha^n$ and $\alpha^{-n}$ don't cluster too much as $n$ grows. Specifically, the gaps between the break points shrink more slowly than any exponential. We'll say $\Sigma$ satisfies the fat gap condition if all its well-cut segments do.

**A global sufficient condition** The fat gap condition, as stated, can be a bit of a pain to check, because it says something about the first return map on every well-cut segment. Fortunately, there's a more globally defined condition which implies the fat gap condition, and is still weak enough to hold for generic translation surfaces.

Following the terminology of [16], let's say a translation surface is $\varphi$-Diophantine if it satisfies the conclusion of part (1) of Theorem 1.1 in [17]. Here, $\varphi\colon [0, \infty) \to (0, \infty)$ is a strongly decreasing function.

**Proposition 3.5.H.** *Suppose $\Sigma$ has no saddle connections, and the vertical flow on $\Sigma \smallsetminus \mathfrak{W}$ is uniquely ergodic. If $\Sigma$ is $(K^{-\lambda t})$-Diophantine for every $\lambda > 0$, then $\Sigma$ satisfies the fat gap condition.*

For any $\lambda > 0$, Proposition 1.2 of [17] can be used to show that if you rotate a translation structure through a full circle, almost all the structures you pass through will be $(K^{-\lambda t})$-Diophantine. Hence, in any rotation family of translation structures, almost all will be $(K^{-\lambda t})$-Diophantine for every $\lambda \in \{\frac{1}{m} : m \in \mathbb{N}_{>0}\}$, and thus for every $\lambda > 0$.

To prove Proposition 3.5.H, we'll need to introduce the notion of a *generalized saddle connection* on a translation surface. This is a geodesic, with respect to the flat metric, from one singularity to another. Equivalently, it's a saddle connection with respect to a rotated translation structure. In [17], Marchese refers to generalized saddle connections simply as saddle connections.

Every translation surface has a canonical complex-valued 1-form $\omega$ that sends horizontal unit vectors to 1 and vertical unit vectors to $i$. In a local translation chart $z\colon \Sigma \to \mathbb{R}^2$, thinking of $\mathbb{R}^2$ as $\mathbb{C}$, this is just the 1-form $dz$. For any generalized saddle connection $\gamma$ on $\Sigma$, let

$$Q_\gamma = \int_\gamma \omega.$$

The real and imaginary parts of $Q_\gamma$ measure the horizontal and vertical distances traveled by $\gamma$.

*Proof of Proposition 3.5.H.* Pick any well-cut segment $Z$ in $\Sigma$. As usual, let $W^+$ and $W^-$ be the places where the backward- and forward-critical leaves pierce $Z$, and let $B^\pm \subset W^\pm$ be the break points of $\alpha^{\mp 1}$. Because $\Sigma$ has no saddle connections, $W^+$ and $W^-$ are disjoint, so the relation $\alpha^\pm$ restricts to a function on $W^\pm$.

For each $w \in W^\pm$ and $n \in \mathbb{N}$, let $\tau_w^{\pm n}$ to be the time at which the vertical flow carries $w$ to $\alpha^{\pm n} w$. Notice that when $\tau_w^n$ is defined, it has the same sign as $n$: positive for $w \in W^+$, and negative for $w \in W^-$. Because the vertical flow on $\Sigma \smallsetminus \mathfrak{W}$ is uniquely ergodic, we can modify the argument after Lemma 1 of [18] to show that

$$\lim_{n \to \pm\infty} \frac{\tau_w^n}{n}$$

converges for all $w \in W^\pm$. The convergence is uniform as $w$ varies over $B^+ \cup B^-$, because $B^\pm$ are finite. We can therefore find a constant $T$ such that $|\tau_b^n| < |n|T$ for all $b \in B^+ \cup B^-$ and $n \in \mathbb{Z}$.

For each $b \in B^\pm$, let $-\rho_b$ be the time at which the vertical flow carries $b$ into its singularity $\mathfrak{b}_b$. Notice that $\rho_b$ is positive for $b \in B^+$ and negative for $b \in B^-$. Let $P$ be the maximum of $|\rho_b|$ over all break points $b$.

For each $n \in \mathbb{N}$, find a pair of points in $W_{\geq K^{-n}}$ that are as close together as possible, so the distance between then is $\mathrm{gap}_{K^{-n}}$. Because they're in $W$, these points can be expressed as $\alpha^k u$ and $\alpha^l v$ for some break points $u, v \in B^+ \cup B^-$. For concreteness, let's say $\alpha^k u$ is to the left of $\alpha^l v$. Because $\alpha^k u$ and $\alpha^l v$ are in $W_{\geq K^{-n}}$, we know that $|k|$ and $|l|$ are at most $n$.

All the break points of $\alpha^{\pm(n+1)}$ are in $W_{\geq K^{-n}}$, so none of them can lie between $\alpha^k u$ and $\alpha^l v$. Thus, as $t$ varies from 0 to $-\tau_u^k$, the vertical flow $\psi^t$

keeps the interval $(\alpha^k u, \alpha^l v)$ in one piece, eventually bringing it to the interval $(u, \alpha^{l-k} v)$. In fact, if we nudge $t$ out to $-\tau_u^k - \rho_u$, the interval still holds together, coming to rest with its left endpoint at $\mathfrak{b}_u$. Similarly, $\psi^t$ keeps the interval $(\alpha^k u, \alpha^l v)$ in one piece as $t$ varies from 0 to $-\tau_v^l - \rho_v$, finally parking it with its right endpoint at $\mathfrak{b}_v$. The vertical flow from times $-(\tau_u^k + \rho_u)$ to $-(\tau_v^l + \rho_v)$ therefore sweeps out a flow box $U_n$ with $\mathfrak{b}_u$ at one corner and $\mathfrak{b}_v$ at the opposite corner.

The diagonal of $U_n$ from $\mathfrak{b}_u$ to $\mathfrak{b}_v$ is a generalized saddle connection, which I'll call $\gamma_n$. Recall that $\gamma_n$ was constructed from a pair of points in $W_{\geq K^{-n}}$ that are as close together as possible. There could be several such pairs, but it won't matter which one we picked. We can find $Q_{\gamma_n}$ by observing that its real and imaginary parts are the width and height of $U_n$, with appropriate signs. The real part is the distance from $\alpha^k u$ to $\alpha^l v$, which we set up to be $\operatorname{gap}_{K^{-n}}$. The imaginary part is $-(\tau_v^l + \rho_v) + (\tau_u^k + \rho_u)$. Letting $\Delta$ be the length of the segment $Z$, we get the bound

$$\begin{aligned} |Q_{\gamma_n}| &\leq |-(\tau_v^l + \rho_v) + (\tau_u^k + \rho_u)| + \operatorname{gap}_{K^{-n}} \\ &\leq |\tau_v^l| + |\tau_u^k| + |\rho_v| + |\rho_u| + \Delta. \end{aligned}$$

Recalling the constants $T$ and $P$ we defined earlier, we can simplify this to

$$\begin{aligned} |Q_{\gamma_n}| &\leq |l|T + |k|T + 2P + \Delta \\ &\leq 2nT + 2P + \Delta. \end{aligned}$$

Now, suppose $\Sigma$ is $(K^{-\lambda t})$-Diophantine for all $\lambda > 0$. This means that, for any $\lambda > 0$, there are only finitely many generalized saddle connections $\gamma$ with

$$|\operatorname{Re} Q_\gamma| < K^{-\lambda |Q_\gamma|}.$$

We want to prove $Z$ satisfies the fat gap condition. In other words, given any $\mu > 0$, we want to show that

$$K^{\mu n} \operatorname{gap}_{K^{-n}}$$

is bounded below by a positive number as $n$ varies in $\mathbb{N}$.

Pick any $\mu > 0$, and let $\lambda = \frac{\mu}{2T}$. Because $\Sigma$ is minimal, $\operatorname{gap}_{K^{-n}}$ goes to zero as $n$ grows, so each generalized saddle connection has only finitely many chances to appear in the sequence $\gamma_n$. Hence, by the Diophantine condition,

$$|\operatorname{Re} Q_{\gamma_n}| < K^{-\lambda |Q_{\gamma_n}|}$$

for only finitely many $n \in \mathbb{N}$. Because $K^{-\lambda t}$ is decreasing, we can infer from our earlier bound on $|Q_{\gamma_n}|$ that $K^{-\lambda(2nT+2P+\Delta)}$ is no greater than $K^{-\lambda |Q_{\gamma_n}|}$, so

$$|\operatorname{Re} Q_{\gamma_n}| < K^{-\lambda(2nT+2P+\Delta)}$$

for only finitely many $n \in \mathbb{N}$. Recalling that $|\operatorname{Re} Q_{\gamma_n}| = \operatorname{gap}_{K^{-n}}$ and rearranging the right-hand side, we see that

$$\operatorname{gap}_{K^{-n}} < K^{-\lambda(2P+\Delta)} K^{-\mu n}$$

for only finitely many $n \in \mathbb{N}$. In other words,

$$K^{\mu n} \operatorname{gap}_{K^{-n}} < K^{-\lambda(2P+\Delta)}$$

for only finitely many $n \in \mathbb{N}$. Since the right-hand side is positive and independent of $n$, it follows immediately that the left-hand side is bounded below by a positive number as $n$ varies in $\mathbb{N}$. □

**Implications for the division metric** Now that we have a practical way to enforce the fat gap condition, let's see how it implies the hypothesis of Corollary 3.3.L, restricting the behavior of Hölder functions on horizontal segments.

**Proposition 3.5.I.** *Let $Z$ be a well-cut segment in $\Sigma$. Suppose $Z$ satisfies the fat gap condition. The gap function on $Z$ then has the property that for any $\lambda > 0$, we can find a positive constant $M$ with $\operatorname{gap}_r \geq Mr^\lambda$.*

*Proof.* Pick any $\lambda > 0$. By the fat gap condition, there's some $\varepsilon > 0$ such that

$$K^{\lambda n} \operatorname{gap}_{K^{-n}} \geq \varepsilon$$

for all large enough $n \in \mathbb{N}$. Equivalently,

$$\operatorname{gap}_{K^{-n}} \geq \varepsilon K^{-\lambda} K^{-\lambda(n-1)}$$

for all large enough $n \in \mathbb{N}$. That means

$$\operatorname{gap}_{K^{-n}} \geq M(K^{-n+1})^\lambda$$

for some $M > 0$ small enough to absorb the constant $\varepsilon K^{-\lambda}$ and deal with the transient behavior of $\operatorname{gap}_{K^{-n}}$ at small $n$. For every $r > 0$ we have $K^{-n} \leq r \leq K^{-n+1}$ for some $n$, so it follows that

$$\operatorname{gap}_r \geq Mr^\lambda$$

for all $r$. □

## 3.6 Divided surfaces as "deflated" hyperbolic surfaces

As I mentioned earlier, divided surfaces can be seen as intermediates between the flat surfaces with translation structures used in abelianization and the hyperbolic surfaces with maximal geodesic laminations used in Bonahon and Dreyer's related construction [6]. Here's a sketch of how this is supposed to work.

First, I should clarify that the correspondence I'm about to describe really involves half-translation surfaces rather than translation surfaces. Readers familiar with prior work on abelianization will find this quite natural, because half-translation surfaces have some claim to being the true home of abelianization. This is discussed in Section 10.1.

Consider a hyperbolic surface carrying a measured maximal geodesic lamination with no leaves that are closed geodesics. The complement of the lamination

is a finite set of ideal triangles. By collapsing those triangles, you can "deflate" the hyperbolic surface to a flat one, with the leaves of the geodesic lamination becoming the vertical leaves of a half-translation structure. The sides of each triangle buckle and fuse into the critical leaves attached to a singularity with cone angle $3\pi$, with the "midpoints" of the sides becoming the singularity itself.

Go... going... gone

The divided version of the half-translation surface, built in essentially the same way as the divided version of a translation surface, is what you might imagine the collapsing hyperbolic surface looks like in the instant before it flattens out. The left and right lanes of the critical leaves are the edges of the triangles, about to fuse together. The medians that separate them are the last vestiges of the interiors of the triangles.

# 4 Warping local systems on divided surfaces

## 4.1 Overview

Many classic geometric constructions involve cutting, shifting, and regluing a local system on a manifold along something akin to a codimension-one submanifold. For example, a Fenchel-Nielsen twist shifts the $\mathrm{PSL}_2\,\mathbb{R}$ local system encoding the hyperbolic structure of a surface along a closed geodesic. A grafting or a cataclysm shifts a $\mathrm{PSL}_2\,\mathbb{R}$ local system along a geodesic lamination on a hyperbolic surface. Higher versions of these processes act on $\mathrm{PSL}_2\,\mathbb{C}$ and $\mathrm{PSL}_n\,\mathbb{R}$ local systems [19][20].

The version of abelianization described in this paper shifts an $\mathrm{SL}_2\,\mathbb{R}$ local system along the critical leaves of a translation surface $\Sigma$. It's most conveniently carried out by pushing the local system up to the divided surface $\overleftrightarrow{\Sigma}$ and warping it along a deviation supported on $\overleftrightarrow{\Sigma}$. The special class of deviations used in this process will be the subject of this section.

Our discussion of warping only makes sense on a locally connected space, so the local connectedness of the divided surface is now playing an important role. The fact that we can warp local systems on the divided surface, like the equivalence of categories of local systems given by Theorem 3.4.E, can be seen as a reflection of the general idea that local systems on the divided surface tend to be well-behaved.

## 4.2 Critical leaves in a flow box

Consider a flow box $U = H \times L$ in $\overset{\leftrightarrow}{\Sigma}$. The median of each critical leaf passes through $U$ at most countably many times, intersecting it in a collection of vertical lines I'll call *dividers*. The dividers are naturally ordered by their positions in $H$. Given two points $y$ and $x$ in $\overset{\cdot\cdot}{U}$, with $y$ sitting to the left of $x$ along $H$, let's write $(y \mid x)^U$ to denote the ordered set of dividers in $U$ that lie between $y$ and $x$.

## 4.3 Deviations defined by jumps, conceptually

Let $\mathcal{F}$ be a locally constant sheaf on $\overset{\leftrightarrow}{\Sigma}$. The $\mathcal{F}$-simple flow boxes form a basis for the topology of $\overset{\leftrightarrow}{\Sigma}$. To specify a *jump* in $\mathcal{F}$, we give for each divider $P$ in a simple flow box $U$ an automorphism $j_P$ of $\mathcal{F}_U$. These automorphisms have to fit together as follows:

- If the basis element $U$ contains the basis element $V$, and the divider $P$ in $U$ contains the divider $Q$ in $V$, the automorphisms $j_Q$ and $j_P$ commute with the restriction morphism $\mathcal{F}_{V \subset U}$.

Consider a simple flow box $U$ and a pair of points $y$ and $x$ in $\overset{\cdot\cdot}{U}$, with $y$ to the left of $x$. In the presence of a jump $j$, the ordered set $(y \mid x)^U$ of dividers between $y$ and $x$ becomes an ordered set of automorphisms of $\mathcal{F}_U$, which are just begging to be composed. There are probably infinitely many of them, so composing them may not make sense, but let's do it anyway, and call the result

$$\delta^U_{yx} = \prod_{P \in (y|x)^U} j_P.$$

For good measure, define $\delta^U_{xy}$ to be the inverse of $\delta^U_{yx}$, so we don't have to worry about checking that $y$ is to the left of $x$ in $U$.

Our notation makes $\delta$ look like a deviation from $\mathcal{F}$ with support $\overset{\cdot\cdot}{\Sigma}$, defined over the basis of $\mathcal{F}$-simple flow boxes. In fact, $\delta$ really will be a deviation as long as the compositions defining it make sense, and behave in the way you'd expect. We'll see this concretely in the next section, where we specialize to the case of jumps in local systems.

## 4.4 Deviations defined by jumps, concretely

Let's say $\mathcal{F}$ is a $G$ local system. In this case, a jump in $\mathcal{F}$ assigns an element of $G$ to each divider. Because $G$ is a linear group, it comes with a topology that can be used to make sense of infinite ordered products, as described in Appendix C. Using the properties of these products, we can show that the construction in the previous section really does produce a deviation $\delta$ from a jump $j$, as long as all the products converge.

### 4.4.1 The restriction property

We want to prove that for any simple flow boxes $V \subset U$ and any two points $y, x \in V \cap \overleftrightarrow{\Sigma}$, the automorphisms $\delta_{yx}^U$ and $\delta_{yx}^V$ commute with the restriction morphism $\mathcal{F}_{V \subset U}$. We can assume, without loss of generality, that $y$ is to the left of $x$. Jumps are required to commute with restriction, so

$$\prod_{P \in (y|x)^U} j_P = \prod_{P \in (y|x)^V} \mathcal{F}_{V \subset U}^{-1} \, j_P \, \mathcal{F}_{V \subset U}.$$

Recalling that conjugation is a topological group automorphism, we apply Proposition C.4.A and get

$$\prod_{P \in (y|x)^U} j_P = \mathcal{F}_{V \subset U}^{-1} \left( \prod_{P \in (y|x)^V} j_P \right) \mathcal{F}_{V \subset U},$$

which is what we wanted to show.

### 4.4.2 The composition property

We want to prove that $\delta_{zy}^U \delta_{yx}^U = \delta_{zx}^U$ for any simple flow box $U$ and any three points $z, y, x \in U \cap \overleftrightarrow{\Sigma}$. If $z, y, x$ happen to be ordered from left to right, we're trying to prove that

$$\left( \prod_{P \in (z|y)^U} j_P \right) \left( \prod_{P \in (y|x)^U} j_P \right) = \prod_{P \in (z|x)^U} j_P.$$

Since, in the notation of Section C, $(z \mid y)^U \sqcup (y \mid x)^U = (z \mid x)^U$, this follows directly from Proposition C.3.A.

Now, suppose the three points are ordered differently. If the ordering from left to right is $y, z, x$, we can use the reasoning above to conclude that $\delta_{yz}^U \delta_{zx}^U = \delta_{yx}^U$, rewrite this as $(\delta_{zy}^U)^{-1} \delta_{zx}^U = \delta_{yx}^U$, and rearrange to get the desired result. The other cases work similarly.

# 5 Uniformity for $\mathrm{SL}_2 \mathbb{R}$ dynamics

## 5.1 The global version

Parallel transport along the vertical flow $\psi \colon \mathbb{R} \times \overleftrightarrow{\Sigma} \to \overleftrightarrow{\Sigma}$ induces a flow $\Psi$ on the stalks of $\mathcal{E}$. At time $t$, the parallel transport flow gives a morphism $\Psi^t \colon \mathcal{E}_x \to \mathcal{E}_{\psi^t x}$ for each $x \in \overleftrightarrow{\Sigma}$. Pick an inner product on the stalks of $\mathcal{E}$ over $\overleftrightarrow{\Sigma}$ which is continuous in the sense that, for any two vectors $u, v \in \mathcal{E}_U$ over an open set $U \subset \overleftrightarrow{\Sigma}$, the inner product of $u$ and $v$ in the stalk $\mathcal{E}_x$ varies continuously as a function of $x \in \overleftrightarrow{\Sigma}$. Because $\overleftrightarrow{\Sigma}$ is compact, it doesn't matter which inner product we pick.

Saying $\mathcal{E}$ is *globally uniform* means that for every $x \in \overset{\leftrightarrow}{\Sigma}$, the dynamics of $\Psi$ split $\mathcal{E}_x$ into a direct sum of two one-dimensional subspaces $\mathcal{E}_x^+$ and $\mathcal{E}_x^-$, called the *forward-stable* and *backward-stable* lines, respectively. These lines are characterized by the following properties:

- The parallel transport map $\Psi^t$ sends $\mathcal{E}_x^\pm$ to $\mathcal{E}_{\psi^t x}^\pm$.
- There is a constant $\Lambda > 0$ such that

$$\lim_{t\to\infty} \frac{1}{t} \log \frac{\|\Psi^{\pm t} v\|}{\|v\|} = -\Lambda \qquad \lim_{t\to\infty} \frac{1}{t} \log \frac{\|\Psi^{\mp t} v\|}{\|v\|} = \Lambda$$

  for all $x \in \overset{\leftrightarrow}{\Sigma}$ and $v \in \mathcal{E}_x^\pm$.
- Each of the limits above converges uniformly. In other words, for any neighborhood $\Omega$ of the right-hand side, making $t$ large enough guarantees that the left-hand side will be in $\Omega$ for all $x$ and $v$.

The constant $\Lambda$ is called the *Lyapunov exponent* of $\mathcal{E}$.

## 5.2 The local version

Consider a simple, well-cut flow box $U = H \times L$ in $\overset{\leftrightarrow}{\Sigma}$, and let $E = \mathcal{E}_U$. Let $\alpha\colon \ddot{H} \to \ddot{H}$ be the first return map discussed in Section 3.5.2. For each $s \in \ddot{H}$, parallel transport along the leaf through $\{s\} \times L$ gives an automorphism $A_s$ of $E$, defined by the commutative square

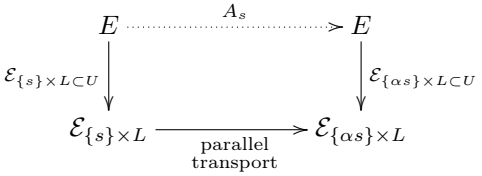


These automorphisms form a cocycle over $\alpha$, which we'll call the *parallel transport cocycle*. Write $A_s^n$ for the parallel transport along $n$ iterations of $\alpha$, starting at $s$. Pick an inner product on $E$ (it doesn't matter which one).

Saying $\mathcal{E}$ is *locally uniform* with respect to $U$ means that for every $s \in \ddot{H}$, the dynamics of $A$ split $E$ into a direct sum of two one-dimensional subspaces $E_s^+$ and $E_s^-$, called the *forward-stable* and *backward-stable* lines, respectively. These lines are characterized by the following properties:

- The parallel transport map $A_s$ sends $E_s^\pm$ to $E_{\alpha s}^\pm$.
- There is a constant $\Lambda > 0$ such that

$$\lim_{n\to\infty} \frac{1}{n} \log \frac{\|A_s^{\pm n} v\|}{\|v\|} = -\Lambda \qquad \lim_{n\to\infty} \frac{1}{n} \log \frac{\|A_s^{\mp n} v\|}{\|v\|} = \Lambda$$

  for all $s \in \ddot{H}$ and $v \in E_s^\pm$.

- Each of the limits above converges uniformly. In other words, for any neighborhood $\Omega$ of the right-hand side, making $n$ large enough guarantees that the left-hand side will be in $\Omega$ for all $s$ and $v$.

## 5.3 The two versions are usually equivalent

Suppose that $\Sigma$ is minimal, and the vertical flow on $\Sigma \smallsetminus \mathfrak{W}$ is uniquely ergodic. Then, for any simple, well-cut flow box $U = H \times L$ in $\overset{\leftrightarrow}{\Sigma}$, global uniformity is equivalent to local uniformity with respect to $U$. That means we can drop the distinction between them, and just talk about *uniformity*.

To see why the two versions of uniformity coincide, pick a horizontal slice $Z = H \times \{\zeta\}$ across $U$. For each $s \in \ddot{H}$, let $\tau_s^n$ be the time it takes for the vertical flow to carry $(s, \zeta)$ back to $Z$ for the $n$th time. Notice that $\psi^{\tau_s^n}(s, \zeta) = (\alpha^n s, \zeta)$. Because the vertical flow on $\overset{\leftrightarrow}{\Sigma}$ is uniquely ergodic, we can imitate the argument after Lemma 1 of [18] to show that

$$\lim_{n\to\infty} \frac{\tau_s^{\pm n}}{\pm n}$$

converges to the same average $\bar{\tau}$ for all $s \in \ddot{H}$.

Because $H$ is full, we can find a continuous function on $\overset{\leftrightarrow}{\Sigma}$ which is one on $Z$ and zero outside of $U$. Using this function, we can pick the continuous inner product on the stalks of $\mathcal{E}$ over $\overset{\leftrightarrow}{\Sigma}$ so that it matches the inner product on $E$ at every point in $Z$. Then, for any $s \in \ddot{H}$ and $v \in \mathcal{E}_{(s,\zeta)}$, we can say $\|A_s^{\pm n} v\| = \|\Psi^{\tau_s^{\pm n}} v\|$. It follows that

$$\begin{aligned}
\lim_{n\to\infty} \frac{1}{n} \log \frac{\|A_s^{\pm n} v\|}{\|v\|} &= \lim_{n\to\infty} \left(\frac{\tau_s^{\pm n}}{\pm n}\right) \left(\frac{\pm 1}{\tau_s^{\pm n}}\right) \log \frac{\|\Psi^{\tau_s^{\pm n}} v\|}{\|v\|} \\
&= \lim_{n\to\infty} \bar{\tau} \left(\frac{\pm 1}{\tau_s^{\pm n}}\right) \log \frac{\|\Psi^{\tau_s^{\pm n}} v\|}{\|v\|} \\
&= \bar{\tau} \lim_{t\to\infty} \frac{1}{t} \log \frac{\|\Psi^{\pm t} v\|}{\|v\|}
\end{aligned}$$

for all $s \in \ddot{H}$ and $v \in \mathcal{E}_{(s,\zeta)}$. Because the vertical flow is minimal, it carries every point in $\overset{\leftrightarrow}{\Sigma}$ through the open set $U$, and thus through $Z$. The identity above therefore shows that $\mathcal{E}$ is globally uniform if and only if it's locally uniform with respect to $U$. The local Lyapunov exponent is $\bar{\tau}$ times the global one. This reflects the fact that the local exponent is unitless, while the global one has units of inverse time.

## 5.4 Extending over medians

Suppose that $\Sigma$ has no saddle connections, and the vertical flow on $\Sigma \smallsetminus \mathfrak{W}$ is uniquely ergodic. Consider a point $w$ on a critical leaf of $\Sigma$. Every neighborhood

of $\hat{w}$ contains $\overleftarrow{w}$ and $\overrightarrow{w}$, so we can take a colimit over neighborhoods $U$ in the diagram

$$\mathcal{E}_{\overleftarrow{w}} \xleftarrow{\mathcal{E}_{\overleftarrow{w}\subset U}} \mathcal{E}_U \xrightarrow{\mathcal{E}_{\overrightarrow{w}\subset U}} \mathcal{E}_{\overrightarrow{w}}$$

to get isomorphisms

$$\mathcal{E}_{\overleftarrow{w}} \longleftarrow \mathcal{E}_{\hat{w}} \longrightarrow \mathcal{E}_{\overrightarrow{w}}$$

identifying the three stalks. I'll refer to all three as $\mathcal{E}_w$, writing $\overleftarrow{v}$, $\hat{v}$, or $\overrightarrow{v}$ when I want to think of a vector $v \in \mathcal{E}_w$ as an element of one or the other.

Let's say $w$ is on a backward-critical leaf. The left and right lanes of this leaf are never separated by the forward vertical flow, so $\Psi^t\overleftarrow{v} = \Psi^t\overrightarrow{v}$ for all $v \in \mathcal{E}_w$ as long as $t \geq 0$. From the continuity of the inner product and the compactness of $\overleftrightarrow{\Sigma}$, you can show that the difference between $\log\|\Psi^t\overleftarrow{v}\|$ and $\log\|\Psi^t\overrightarrow{v}\|$ stays bounded as $t$ varies over $[0,\infty)$.[6] Thus, if $\mathcal{E}$ is uniform, the forward-stable lines $\mathcal{E}^+_{\overleftarrow{w}}$ and $\mathcal{E}^+_{\overrightarrow{w}}$ are the same, and we can refer to both as $\mathcal{E}^+_w$. The backward-stable lines at $\overleftarrow{w}$ and $\overrightarrow{w}$, on the other hand, are typically different. If $w$ is on a forward-critical leaf instead of a backward-critical one, we can use the same reasoning in the other direction to define $\mathcal{E}^-_w$.

Working in a well-cut flow box $H \times L \subset \overleftrightarrow{\Sigma}$, it will be useful to describe our extension of $\Psi^t$ in terms of the first return map. As we did in Section 3.5.2, let $W \subset \pi H$ be the positions of the critical leaves, and partition it into the backward- and forward-critical sets $W^+$ and $W^-$. By our previous reasoning, $A_{\overleftarrow{w}} = A_{\overrightarrow{w}}$ for all $w \in W^+$. Defining $A_{\hat{w}}$ to be equal to both, we can extend the forward parallel transport cocycle $A$ over the medians of all backward-critical points. If $\mathcal{E}$ is uniform, the forward-stable lines $E^+_{\overleftarrow{w}}$ and $E^+_{\overrightarrow{w}}$ of $E = \mathcal{E}_U$ match, so we can define $E^+_{\hat{w}}$ to be equal to both. The backward cocycle $A^{-1}$ extends over the medians of all forward-critical points in the same way, allowing us to define $E^-_{\hat{w}}$ for all $w \in W^-$.

## 5.5 Constructing uniform local systems

Our abelianization process can only be carried out when $\mathcal{E}$ is uniform, so it will be nice to have a way of constructing uniform $\mathrm{SL}_2\,\mathbb{R}$ local systems on $\overleftrightarrow{\Sigma}$. We might as well assume that $\Sigma$ is minimal and the vertical flow on $\Sigma \smallsetminus \mathfrak{W}$ is uniquely ergodic, since we'll need those conditions for abelianization anyway. The construction of uniform local systems can then be done locally, in a well-cut flow box $H \times L \subset \overleftrightarrow{\Sigma}$.

Recall that $\alpha\colon \dot{\bar{H}} \to \dot{\bar{H}}$ is the fractured version of an interval exchange. The parallel transport cocycle over $\alpha$ is constant on each of the exchanged

[6] Here's a sketch of the proof. Cover $\overleftrightarrow{\Sigma}$ with a finite collection $\mathcal{U}$ of simple open sets. For each $U \in \mathcal{U}$, the closure of $\dot{\bar{U}}$ in $\overleftrightarrow{\Sigma}$ is compact. Hence, the difference between $\log\|v\|_x$ and $\log\|v\|_y$ stays bounded as $x$ and $y$ vary over $\dot{\bar{U}}$ and $v$ varies over $\mathcal{E}_U$. In particular, the difference between $\log\|v\|_{\overleftarrow{w}}$ and $\log\|v\|_{\overrightarrow{w}}$ is bounded over all $w \in \mathfrak{W} \cap \pi U$. Now, just combine the bounds over all $U \in \mathcal{U}$.

intervals. I'll call this kind of cocycle an *interval cocycle*. A local system on $\overleftrightarrow{\Sigma}$ is determined up to isomorphism by the parallel transport cocycle it induces over $\alpha$. Conversely, any interval cocycle over $\alpha$ is the parallel transport cocycle of some local system on $\overleftrightarrow{\Sigma}$.

To get a sense of why the claims above are true, first recall that Theorem 3.4.E lets us pass from local systems on $\overleftrightarrow{\Sigma}$ to ones on $\Sigma \smallsetminus \mathfrak{B}$. Because $\Sigma$ is minimal, we can express $\Sigma \smallsetminus \mathfrak{B}$ as a suspension of the first return relation on $\iota^{-1}H$ [15]; you can see roughly what that means by looking at the proof of Proposition 3.5.G. Any local system on $\Sigma \smallsetminus \mathfrak{B}$ can be trivialized over the flow boxes that make up the suspension, with transition morphisms given by the parallel transport cocycle. Conversely, the transition morphisms given by an interval cocycle can be used to construct a local system over $\Sigma \smallsetminus \mathfrak{B}$.

Now all we need is a way of constructing uniform $\mathrm{SL}_2\,\mathbb{R}$ cocycles over $\alpha\colon \ddot{H} \to \ddot{H}$. For convenience, let's have our cocycles act on $\mathbb{R}^2$ with the standard volume form. Let $\mathcal{A}$ be the intersections of the intervals exchanged by $\alpha$ and the ones exchanged by $\alpha^{-1}$. Consider a pair of functions $u, v\colon \mathcal{A} \to \mathbb{R}^2$ that assign a basis for $\mathbb{R}^2$ to every interval in $\mathcal{A}$. We can think of $u$ and $v$ as continuous functions $\ddot{H} \to \mathbb{R}^2$ by composing them with the function $\ddot{H} \to \mathcal{A}$ that sends each point to its interval. An interval cocycle over $\alpha$ is *positive* with respect to $u, v$ if at every $x \in \ddot{H}$ it maps $u_x$ and $v_x$ into the interior of the cone generated by $u_{\alpha x}$ and $v_{\alpha x}$. A cocycle is *eventually positive* if its $n$th iteration is positive, for some $n \in \mathbb{N}$. Theorems 3 and 4 of [7] show that an interval cocycle over $\alpha$ is uniform if and only if it's eventually positive with respect to some basis $u, v\colon \mathcal{A} \to \mathbb{R}^2$.

Using this result, you can easily construct a lot of uniform $\mathrm{SL}_2\,\mathbb{R}$ local systems on $\Sigma \smallsetminus \mathfrak{B}$: just write down interval cocycles whose matrix entries are all positive with respect to the standard basis for $\mathbb{R}^2$. You can also see that uniform cocycles form an open set in the space of all interval cocycles, which is a product of copies of $\mathrm{SL}_2\,\mathbb{R}$.

# 6 Abelianization in principle

## 6.1 Overview

Now that we have the tools we need, we can turn again to our goal of extending abelianization to surfaces without punctures. At this point, it will be useful to take a closer look at the original description of abelianization, which is implicit in [5, §10], but first appears explicitly in [4, §4].

Our review will be simplified in two important ways. First, we'll restrict our attention to $\mathrm{SL}_2\,\mathbb{C}$ local systems, though Gaiotto, Hollands, Moore, and Neitzke show how to abelianize special linear local systems of any rank. Second, we'll only talk about abelianization using translation structures, leaving aside Gaiotto, Moore, and Neitzke's more powerful and general *spectral networks*.

### 6.1.1 Setting

Our review takes place on a translation surface $\Sigma'$ which is compact except for a finite set of *punctures*. A puncture is a region homeomorphic to a punctured disk, with a translation structure taken from a certain family of shapes. This definition of a puncture is analogous to our earlier definition of a singularity. For simplicity, let's consider a translation surface whose punctures all have the simplest shape: a half-infinite cylinder. A complete list of puncture shapes, and an explanation of where they come from, can be found in Appendix E.

Let's assume $\Sigma'$ has no saddle connections. In this case, if you follow a vertical leaf in some direction, your fate is easy to describe. If you're following a critical leaf in the critical direction, you will by definition fall into a singularity after a finite amount of time. Otherwise, you'll end up falling forever into a puncture; in our case, that means spiraling down the long end of a half-infinite cylinder.

Every non-critical leaf on $\Sigma'$ is thus associated with two punctures, not necessarily distinct: the punctures its ends spiral into. If you remove the critical leaves $\mathfrak{W}$, the surface $\Sigma'$ falls apart into infinite vertical strips of leaves that all go into the same punctures. Each strip is bounded by four critical leaves, joined at two singularities, as illustrated below. Gaiotto, Hollands, Moore, and Neitzke compactify the surface by filling in the punctures, so the closure of each strip becomes a quadrilateral with singularities as two of its vertices and punctures as the other two.

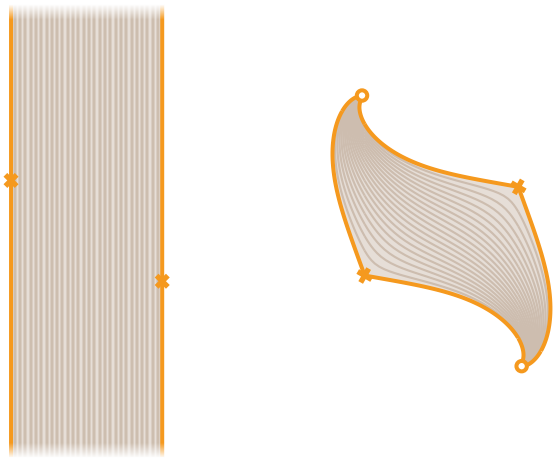

Strip Compactified strip

### 6.1.2 Framings

To abelianize an $\mathrm{SL}_2\,\mathbb{C}$ local system $\mathcal{E}$ on $\Sigma'$, we first need to give it a bit of extra structure, called a *framing* (or *flag data*, in [5]). In the abstract settings of [4] and [5], a framing is specified geometrically by giving a projectively flat section of $\mathcal{E}$

on a neighborhood of each puncture (a framing can also be specified analytically, as we'll recall in Section 6.1.4). For reasons that will become apparent later, I'll refer to the sections given by the framing as *stable lines*. Framings always exist, because an operator on a finite-dimensional complex vector space always has at least one eigenvector.

The framing gives a pair of lines in every stalk of $\mathcal{E}$ over a non-critical leaf. One, which I'll call the *forward-stable* line, is gotten by following the forward vertical flow until you fall into a puncture, grabbing the stable line, and carrying it back by parallel transport. The *backward-stable* line is gotten in the same way by following the backward vertical flow. The forward- and backward-stable lines fit together into sections of $\mathcal{E}$ over every strip of $\Sigma' \smallsetminus \mathfrak{W}$.

The framing also gives a line in every stalk of $\mathcal{E}$ over a critical leaf—the stable line from the puncture the critical leaf falls into. This line matches the backward- or forward-stable lines in the strips on either side, depending on whether the leaf is forward- or backward-critical. Hence, as you cross the boundary between two strips, one of the stable lines stays fixed, although the other can change.

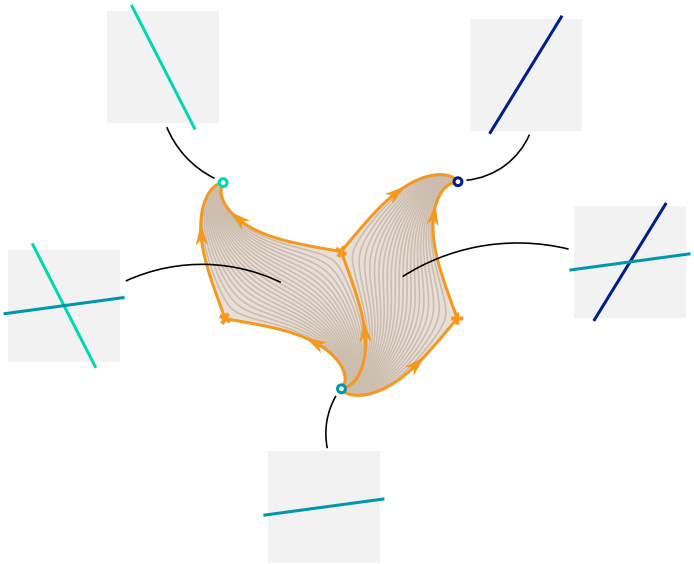

The stable lines over two adjacent strips

### 6.1.3 Abelianization

For a generic local system, the forward and backward-stable lines over each strip are complementary, splitting $\mathcal{E}$ into a direct sum of $\mathbb{C}^\times$ local systems over $\Sigma' \smallsetminus \mathfrak{W}$. The changes in the stable lines at the boundaries between strips generally prevent this splitting from extending over all of $\Sigma'$. At each critical

leaf, however, there's a unique element of $\mathrm{SL}_2\,\mathbb{C}$ that sends the stable lines on one side to the stable lines on the other, acting by the identity on the line that stays fixed.

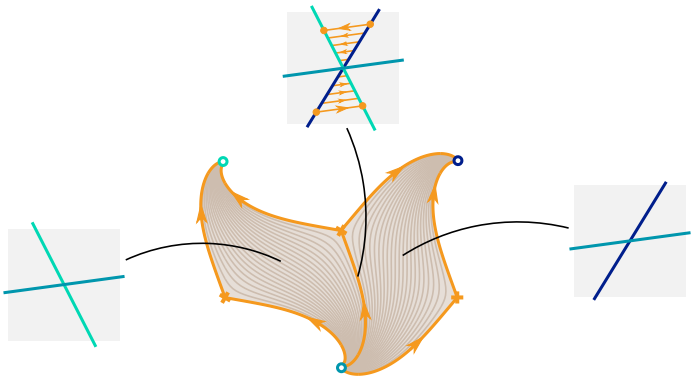

By cutting $\mathcal{E}$ along the critical leaves, shifting it by this automorphism, and gluing it back together, we can match up the stable lines across the boundaries of the strips, so the splitting they give becomes global.

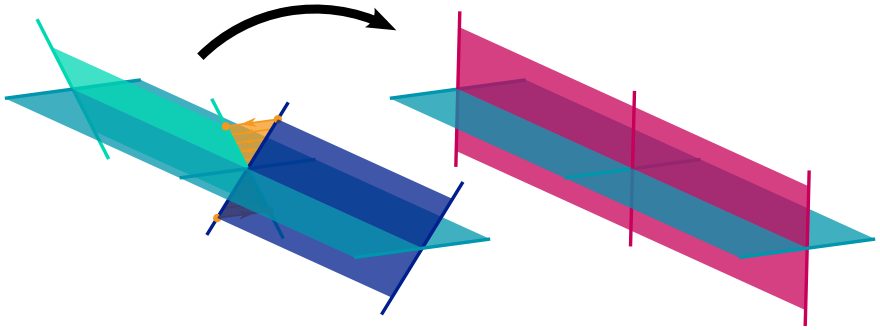

Cutting and gluing to match up the stable lines

That's abelianization.

### 6.1.4 Abelianization without punctures

If we want to carry out the process above on a surface without punctures, there are two questions we have to settle. One is conceptual: what should play the role of the framing? If you've read Section 5, our suggestive terminology may have already told you the answer. When $\mathcal{E}$ is uniform, it comes with complementary forward- and backward-stable lines over every non-critical leaf, which can be used as the forward- and backward-stable lines in the process above. The discussion in Section 5.4 amounts to saying that one of the stable lines stays

fixed when you cross a critical leaf, so we can define the automorphisms over the critical leaves just as we did before.

This approach feels different than the one in [5], because the stable lines of a uniform local system are a property of the local system, rather than additional data. It's very reminiscent, however, of the more concrete approach taken in [21], where the framing comes from a specially chosen inner product on the stalks of the local system.[7] The stable line at a puncture consists of the sections that shrink as you fall in, just as the stable lines in our approach consist of the sections that shrink as you follow the vertical flow.

The stable lines in [5] and [21] live at the punctures in the surface, which lift to points on the boundary of the universal cover. Similarly, the stable lines in our approach can be seen as living on the boundary of the universal cover, at the points where the lifts of the vertical leaves begin and end. This point of view is central to Bonahon and Dreyer's version of abelianization [6], but it plays little role in ours.

The other question is just a technical difficulty. On a compact translation surface with no saddle connections, every vertical leaf is dense, so how do we think about shifting $\mathcal{E}$ along the automorphisms over the critical leaves? How do we know the process is well-defined? How do we know the resulting local system actually splits into a direct sum of $\mathbb{C}^\times$ local systems? The three parts of this question are answered in Sections 6.3, 7, and 8, respectively.

## 6.2 Running assumptions

The compact translation surface $\Sigma$ introduced in Section 1.4.1 will, of course, stay with us. We'll discuss the abelianization of a fixed $\mathrm{SL}_2\,\mathbb{R}$ local system $\mathcal{E}$ on $\overleftrightarrow{\Sigma}$. Abelianization will yield a new $\mathrm{SL}_2\,\mathbb{R}$ local system $\mathcal{F}$ and a stalkwise isomorphism $\Upsilon\colon \mathcal{E} \to \mathcal{F}$, supported on $\overleftrightarrow{\Sigma}$. As discussed in Section 1.2.3, our results should generalize from $\mathrm{SL}_2\,\mathbb{R}$ to $\mathrm{SL}_2\,\mathbb{C}$ without too much trouble.

From now on, we'll assume the following things about $\Sigma$:

- The surface $\Sigma$ has no saddle connections.
- The vertical flow on $\Sigma \smallsetminus \mathfrak{W}$ is uniquely ergodic.
- The surface $\Sigma$ satisfies the fat gap condition of Section 3.5.5.

We'll also make one crucial assumption about $\mathcal{E}$:

- The local system $\mathcal{E}$ is uniform.

If you happen to have picked a surface $\Sigma$ that doesn't satisfy the required conditions, have no fear. The remarks in Sections 3.2.1 and 3.5.5 show that you can fix this problem with an arbitrarily small rotation of the translation structure on $\Sigma$. In fact, if you rotate the translation structure on $\Sigma$ through a full circle, all but a measure-zero subset of the structures you pass through will

[7] That inner product is the celebrated *harmonic metric* from the theory of Hitchin systems. See Sections 6.5 and 13.2 of [21] for more information.

satisfy all the conditions needed. Once you've fixed a good translation structure on $\Sigma$, you can use the results of Section 5.5 to find lots of uniform local systems $\mathcal{E}$.

## 6.3 The slithering jump

### 6.3.1 Overview

Working on $\overset{\leftrightarrow}{\Sigma}$ gives us a convenient way to talk about the stable lines on either side of a critical leaf: using the identification in Section 5.4, we can compare the stable lines over the left and right lanes. The automorphisms that match up the stable lines across the median can be encoded as a jump in the local system $\mathcal{E}$, as defined in Section 4. This jump contains essentially the same information as *slithering maps* defined by Bonahon and Dreyer in [6], so we'll call it the *slithering jump*. We abelianize $\mathcal{E}$ by warping it along the deviation defined by the slithering jump. More explicitly, we abelianize $\mathcal{E}$ by carrying out the following steps:

1. Compute the slithering jump, using the formulas in Section 6.3.3.
2. Turn the slithering jump into a deviation, as described in Section 4.3.
3. Warp $\mathcal{E}$ along the deviation, as described in Section 2.5.

We'll prove in Sections 7 and 8 that these instructions have the desired effect, as long as the conditions in Section 6.2 are satisfied.

### 6.3.2 Definition

Consider a point $w$ on a backward-critical leaf of $\Sigma$. Because $\mathcal{E}_w$ is two-dimensional, and $\mathrm{SL}_2\,\mathbb{R}$ is the group of volume-preserving linear maps, there's a unique automorphism $s_w$ of $\mathcal{E}_w$ that sends $\mathcal{E}^-_{\overrightarrow{w}}$ to $\mathcal{E}^-_{\overleftarrow{w}}$ and is the identity on $\mathcal{E}^+_w$. This induces an automorphism of $\mathcal{E}_U$ for any simple flow box $U$ containing $\hat{w}$. If $w$ is on a forward-critical leaf instead of a backward-critical one, we can define $s_w$ in the same way, with the roles of the backward-stable and forward-stable lines reversed.

Given a divider $P$ in a simple flow box $U$, it's not hard to see that $s_w$ induces the same automorphism of $\mathcal{E}_U$ for every $w \in P$. Call this automorphism $s_P$. As $P$ varies over all dividers in all simple flow boxes, the automorphisms $s_P$ fit together into a jump $s$ in the local system $\mathcal{E}$—the *slithering jump*.

### 6.3.3 Formulas

Assuming, for convenience, that $w$ is on a backward-critical leaf, we can get an explicit expression for $s_w$ by choosing representatives

$$u' \in \mathcal{E}^-_{\overleftarrow{w}} \qquad\qquad v \in \mathcal{E}^+_w \qquad\qquad u \in \mathcal{E}^-_{\overrightarrow{w}}.$$

Observe that $\{v, u'\}$ and $\{v, u\}$ are ordered bases for $\mathcal{E}_w$. By rescaling $u'$ and $u$, we can ensure that both ordered bases have unit volume. The transformation $s_w$ is then given by[8]

$$v \mapsto v \qquad\qquad u \mapsto u'.$$

A quick calculation with the volume form $D$ gives the relation

$$u' = u + D(u', u)\, v,$$

revealing that $s_w$ is a shear transformation whose off-diagonal part $s_w - 1$ is given by

$$v \mapsto 0 \qquad\qquad u \mapsto D(u', u)v.$$

When $w$ is on a forward-critical leaf, similar expressions can be obtained.

### 6.3.4 Flow invariance

Suppose $w$ is on a backward-critical leaf. Because $s_w$ is defined in terms of the stable lines $\mathcal{E}^{\pm}_w$ and the volume form on $\mathcal{E}_w$, which are all invariant under the vertical flow, the diagram

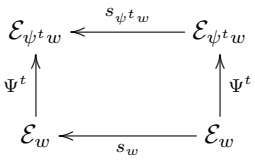


commutes for all positive times $t$. If $w$ is on a forward-critical leaf, the same is true for all negative times.

This flow invariance property is not unique to the slithering jump. In fact, it holds for all jumps, as a direct consequence of the defining consistency condition. Its introduction has been delayed until now only for convenience.

# 7 Abelianization converges

## 7.1 Overview

To show that the slithering jump defines a deviation $\sigma$, as discussed in Section 4, we need to verify that the infinite product defining the automorphism $\sigma^U_{yx}$ converges for every simple flow box $U \subset \overleftrightarrow{\Sigma}$ and every pair of points $y, x \in \dot{\bar{U}}$.

[8] Given a pre-existing basis for $\mathcal{E}_w$, you can find $s_w$ by solving the matrix equation

$$\left[\begin{array}{c|c} v & u' \end{array}\right] = s_w \left[\begin{array}{c|c} v & u \end{array}\right],$$

which I have found convenient for numerical work.

Because jumps commute with restrictions, it's enough to check for convergence on a set of simple flow boxes that cover $\overleftrightarrow{\Sigma}$. We'll use the simple, well-cut flow boxes for this purpose.

Consider a simple, well-cut flow box $U = H \times L$ in $\overleftrightarrow{\Sigma}$, and let $E = \mathcal{E}_U$. Pick an inner product on $E$, so we can use the asymptotic growth conditions given by the uniformity of $\mathcal{E}$. As usual, let $W \subset \pi H$ be the positions of the critical leaves, recalling that $\ddot{H} = H \smallsetminus \iota W$. Label each divider $\{\hat{w}\} \times L$ in $U$ by the point $w \in W$ it sits above. As we did in Section 3.5.2, let $B^+ \subset W^+$ and $B^- \subset W^-$ be the break points of $\alpha^{-1}$ and $\alpha$, respectively.

Throughout this section and the next, for positive functions $f$ and $g$, we'll write $f \lesssim g$ to say that $f$ is bounded by a constant multiple of $g$.

## 7.2 Bounding the jump

For any $b \in B^+$, we can choose representatives

$$u' \in E^-_{\overleftarrow{b}} \qquad\qquad v \in E^+_{\hat{b}} \qquad\qquad u \in E^-_{\overrightarrow{b}}$$

for which $D(v, u')$ and $D(v, u)$ are 1 and conclude that $s_b - 1$ is given by

$$v \mapsto 0 \qquad\qquad u \mapsto D(u', u)\, v,$$

as described in the previous section. Applying the flow invariance of jumps, we see that $s_{\alpha^n b} - 1$ is given by

$$A^n_{\hat{b}} v \mapsto 0 \qquad\qquad A^n_{\hat{b}} u \mapsto D(u', u)\, A^n_{\hat{b}} v$$

for all positive $n$.

We see from the formula above that $D(u', u)\, A^n_{\hat{b}} v$ spans the image of $s_{\alpha^n b} - 1$. We also learn that the shortest vector $s_{\alpha^n b} - 1$ sends to $D(u', u)\, A^n_{\hat{b}} v$ is the orthogonal projection of $A^n_{\hat{b}} u$ onto the orthogonal complement of $A^n_{\hat{b}} v$. From this, we can calculate the operator norm of $s_{\alpha^n b} - 1$:

$$\| s_{\alpha^n b} - 1 \| = \frac{|D(u', u)|}{d_\angle(A^n_{\hat{b}} v, A^n_{\hat{b}} u)} \frac{\| A^n_{\hat{b}} v \|}{\| A^n_{\hat{b}} u \|},$$

where $d_\angle$ is the function that takes two nonzero vectors in $E$ and measures the sine of the angle between the lines they span. Rearranging a bit, we get

$$\| s_{\alpha^n b} - 1 \| = \frac{|D(u', u)|}{d_\angle(A^n_{\hat{b}} v, A^n_{\hat{b}} u)} \frac{\| v \|}{\| u \|} \cdot \frac{\| A^n_{\hat{b}} v \|}{\| v \|} \cdot \frac{\| u \|}{\| A^n_{\hat{b}} u \|}.$$

The argument at the end of Section 8.2.3 shows that $d_\angle(A^n_{\hat{b}} v, A^n_{\hat{b}} u)$ is bounded away from zero. That and the uniformity of $\mathcal{E}$ give us the limit

$$\lim_{n \to \infty} \frac{1}{n} \left[ \log \frac{|D(u', u)|}{d_\angle(A^n_{\hat{b}} v, A^n_{\hat{b}} u)} \frac{\| v \|}{\| u \|} + \log \frac{\| A^n_{\hat{b}} v \|}{\| v \|} - \log \frac{\| A^n_{\hat{b}} u \|}{\| u \|} \right] = -2\Lambda,$$

demonstrating that

$$\lim_{n\to\infty} \frac{1}{n} \log \|s_{\alpha^n b} - 1\| = -2\Lambda$$

for all $b \in B^+$.

Applying the same reasoning in the other direction, we see more generally that

$$\lim_{n\to\infty} \frac{1}{n} \log \|s_{\alpha^{\pm n} b} - 1\| = -2\Lambda$$

for all $b \in B^\pm$, implying that

$$\|s_{\alpha^{\pm n} b} - 1\| \lesssim K^{-n}$$

for any $K < e^{2\Lambda}$. Because there are only finitely many break points in each direction, the constant multiple implied by this expression can be made uniform across all $b \in B^\pm$.

## 7.3 Showing its product converges

Recall that $\sigma$ is the deviation we're hoping will be defined by the slithering jump. Pick any two points $y, x \in \ddot{H} \times L$. Since we're labeling the dividers in $U$ by points of $W$, we can think of $(y \mid x)^U$ as a subset of $W$, and write

$$\sigma_{yx}^U = \prod_{w \in (y|x)^U} s_w.$$

Proposition C.6.A in Appendix C tells us that this product converges if the sum

$$C_{yx} = \sum_{w \in (y|x)^U} \|s_w - 1\|$$

does. (I've given the sum a name because its value, as a function of $y$ and $x$, will be useful to us later.)

Let's say every point in $(y \mid x)^U$ takes at least $n$ iterations of $\alpha$ or $\alpha^{-1}$ to hit a break point. Then the set

$$\bigcup_{m \geq n} \Big[ \{\alpha^m b : b \in B^+\} \cup \{\alpha^{-m} b : b \in B^-\} \Big]$$

contains all the points in $(y \mid x)^U$, so

$$C_{yx} \leq \sum_{m \geq n} \left[ \sum_{b \in B^+} \|s_{\alpha^m b} - 1\| + \sum_{b \in B^-} \|s_{\alpha^{-m} b} - 1\| \right].$$

Choosing some $K \in (1, e^{2\Lambda})$ and applying the bound from the previous section, we see that

$$\begin{aligned} C_{yx} &\lesssim \sum_{m \geq n} \left[ \sum_{b \in B^+} K^{-m} + \sum_{b \in B^-} K^{-m} \right] \\ &\lesssim \sum_{m \geq n} K^{-m}. \end{aligned}$$

Hence, the sum defining $C_{yx}$ converges.

Summing the geometric series, we learn that $C_{yx} \lesssim K^{-n}$. But $n$ is the grade of the highest point in $(y \mid x)^U$, so $K^{-n}$ is the distance between $y$ and $x$ in the division metric! We've proven that $C_{yx} \lesssim d(y, x)$. The constant multiple implied by this expression does not depend on $y$ and $x$.

# 8 Abelianization delivers

## 8.1 Overview

Now we know the slithering jump defines a deviation $\sigma$, so we can warp $\mathcal{E}$ along this deviation to produce a new local system $\mathcal{F}$ and a stalkwise isomorphism $\Upsilon\colon \mathcal{E} \to \mathcal{F}$, supported on $\dot{\dot{\Sigma}}$. By design, $\Upsilon$ matches up the stable lines of $\mathcal{E}$ across the medians of $\overleftrightarrow{\Sigma}$, sending corresponding stable lines in $\mathcal{E}_{\overleftarrow{w}}$ and $\mathcal{E}_{\overrightarrow{w}}$ to the same line in $\mathcal{F}_{\hat{w}}$ for all $w \in \mathfrak{W}$. To show that $\mathcal{F}$ splits into a direct sum of two $\mathbb{R}^\times$ bundles, we need to prove that $\Upsilon$ matches up the stable lines on larger scales. For any simple flow box $U \subset \overleftrightarrow{\Sigma}$, we have to show that $\Upsilon$ sends the corresponding stable lines in all the stalks of $\mathcal{E}$ over $\dot{\dot{U}}$ to the stalk restrictions of a single line in $\mathcal{F}_U$. Because of the way deviations restrict, it's enough to prove the desired result on a set of simple flow boxes that cover $\overleftrightarrow{\Sigma}$, just like in Section 7. Once again, we'll use the simple, well-cut flow boxes for this purpose.

We'll keep all the notation from Section 7, and add to it the shorthand $F = \mathcal{F}_U$. To make the geometry facts from Appendix D available, scale the inner product on $E$ so that the unit square has unit volume. To make the results from Section 3.3.5 available, give $\dot{\dot{H}}$ the division metric with steepness $K$. Make the steepness less than $e^{2\Lambda}$, so we can use it as the parameter $K$ that appears in Section 7.

The argument we're about to do is somewhat technical, so let's first recall how it works over a punctured surface $\Sigma'$, where it's so straightforward that we barely mentioned it earlier. Think of the stable lines of $\mathcal{E}$ as lines in $E$ parameterized by the points of $\dot{\dot{U}}$, and think of their images in $\mathcal{F}$ under $\Upsilon$ as lines in $F$. The stable lines are constant in $E$ away from the critical leaves of $\Sigma'$, and the slithering jump only disturbs $\mathcal{E}$ at the critical leaves. Hence, the images under $\Upsilon$ of the stable lines are constant in $F$ away from the critical leaves. By design, the images of the stable lines under $\Upsilon$ are also constant across the critical leaves, so they must be constant everywhere. In other words, $\Upsilon$ matches up the images of the stable lines in $F$ all across $\dot{\dot{U}}$.

On the unpunctured surface $\Sigma$, it's hard to get away from the critical leaves, which fill the surface densely. The property of being constant away from the critical leaves thus has no clear meaning, and we'll have to replace it with something else if we want to reproduce the argument above. Corollary 3.3.L and the discussion at the end of Section 3.5.5 suggest that in our scenario, Hölder continuity might be a viable substitute. Most of the work below is concerned with proving that the stable lines and their images under $\Upsilon$ vary Hölder continuously over $\dot{\dot{U}}$.

## 8.2 The stable distributions are Hölder

Let's say the distance between two lines in $E$ is the sine of the angle between them. This puts a metric on the projective space $\mathbf{P}E$, which I'll call the *sine metric*. I'll write $d_\angle(u, v)$ to mean the distance in $\mathbf{P}E$ between the lines generated by $u, v \in E$.

Let's collect all the forward- and backward-stable lines over $\ddot{H}$ into a pair of functions $E^\pm \colon \ddot{H} \to \mathbf{P}E$, the *stable distributions*. Our regularity conditions on $\Sigma$ and $\mathcal{E}$ guarantee that $E^\pm$ are Hölder. We can see this by applying a recent theorem of Araújo, Bufetov, and Filip to the parallel transport cocycle over $\ddot{H}$ [22, Theorem A]. Once we check that its conditions are satisfied, the theorem will tell us that $E^\pm$ are Hölder on any appropriate *regular block* in $\ddot{H}$. We'll then show that all of $\ddot{H}$ is an appropriate regular block.

### 8.2.1 The conditions of the theorem are satisfied

As we discussed in Section 3.5.2, the first return map on $\ddot{H}$ is the fractured version of an interval exchange on $\pi H$, with break points $\iota B^-$. Removing the dividers over the break points from $U$ turns it into a disjoint union of full flow boxes $H_1 \times L, \dots, H_k \times L$. Each of these flow boxes stays in one piece as the vertical flow carries it around $\overleftrightarrow{\Sigma}$ and back to $U$. As a result, the parallel transport cocycle $A \colon \ddot{H} \to \operatorname{Aut} E$ is constant on each of the intervals $\ddot{H}_1, \dots, \ddot{H}_k$, which form an open partition of $\ddot{H}$.

It immediately follows that $\log \|A\|$ and $\log \|A^{-1}\|$ are integrable with respect to the ergodic probability measure on $\ddot{H}$. Because the intervals $\ddot{H}_1, \dots, \ddot{H}_k$ are full, distances between points in different intervals are bounded away from zero, so the fact that $A$ is constant on each interval also implies that $A$ is Hölder.

We saw in Section 3.3.5 that the fractured version of an interval exchange is Lipschitz with respect to the division metric it defines. The first return map on $\ddot{H}$ is therefore Lipschitz.

### 8.2.2 The stable lines are Hölder on appropriate regular blocks

Regular blocks are parameterized by two real numbers $\varepsilon > 0$ and $\ell > 1$. Araújo, Bufetov, and Filip take $\ell$, for convenience, to be an integer, and the same could be done here as well. Their proof of Theorem A, found in [22, §3.2], shows that $E^\pm$ are Hölder on any regular block for which $\varepsilon$ is small enough and $\ell$ is large enough. Referring back to [22, §2.2], where the relevant thresholds are defined, we see that $\varepsilon$ must be smaller than a tenth of the minimum gap between the Lyapunov exponents of $A$, and $\ell$ must be large enough for the regular block to have positive measure.

In our case, the Lyapunov exponents of $A$ are $\pm\Lambda$, so the minimum gap is $2\Lambda$. In the next section, we'll choose $\varepsilon$ to be less than $\frac{1}{5}\Lambda$, and we'll make $\ell$ large enough for the resulting regular block to be the whole of $\ddot{H}$, which of course has measure one.

### 8.2.3 The whole interval is an appropriate regular block

In our context, the *regular block* with parameters $\ell \in (1, \infty)$ and $\varepsilon \in (0, \frac{1}{5}\Lambda)$ consists of the points $h \in \ddot{H}$ at which

$$\begin{aligned} \tfrac{1}{\ell} e^{(-\Lambda-\varepsilon)n-\varepsilon|m|}\|v\| \le \|A^{\pm n}_{\alpha^m h} v\| &\le \ell e^{(-\Lambda+\varepsilon)n+\varepsilon|m|}\|v\| \\ \tfrac{1}{\ell} e^{(\Lambda-\varepsilon)n-\varepsilon|m|}\|v\| \le \|A^{\mp n}_{\alpha^m h} v\| &\le \ell e^{(\Lambda+\varepsilon)n+\varepsilon|m|}\|v\| \end{aligned}$$

for all $m \in \mathbb{Z}$, $v \in E^{\pm}_{\alpha^m h}$, and $n \ge 0$, and

$$\sqrt{1 - d_{\angle}(E^{+}_{\alpha^n h}, E^{-}_{\alpha^n h})^2} \le 1 - \tfrac{1}{\ell} e^{-\varepsilon|n|}$$

for all $n \in \mathbb{Z}$ [22, §2.2]. We'll see that, given any $\varepsilon$, we can choose $\ell$ large enough to make both conditions hold for all $h \in \ddot{H}$.

To see that the norm condition is satisfied everywhere for large enough $\ell$, first note that it's enough to show that

$$\begin{aligned} \tfrac{1}{\ell} e^{(-\Lambda-\varepsilon)n}\|v\| \le \|A^{\pm n}_{\alpha^m h} v\| &\le \ell e^{(-\Lambda+\varepsilon)n}\|v\| \\ \tfrac{1}{\ell} e^{(\Lambda-\varepsilon)n}\|v\| \le \|A^{\mp n}_{\alpha^m h} v\| &\le \ell e^{(\Lambda+\varepsilon)n}\|v\| \end{aligned}$$

for all $m \in \mathbb{Z}$, $v \in E^{\pm}_{\alpha^m h}$, and $n \ge 0$. Now $m$ plays no role other than to move our starting point, so really we only need to show that

$$\begin{aligned} \tfrac{1}{\ell} e^{(-\Lambda-\varepsilon)n}\|v\| \le \|A^{\pm n}_{h} v\| &\le \ell e^{(-\Lambda+\varepsilon)n}\|v\| \\ \tfrac{1}{\ell} e^{(\Lambda-\varepsilon)n}\|v\| \le \|A^{\mp n}_{h} v\| &\le \ell e^{(\Lambda+\varepsilon)n}\|v\| \end{aligned}$$

for all $v \in E^{\pm}_{h}$ and $n \ge 0$. Equivalently,

$$\begin{aligned} -\Lambda - \left(\varepsilon + \frac{\log \ell}{n}\right) \le \frac{1}{n} \log \frac{\|A^{\pm n}_h v\|}{\|v\|} &\le -\Lambda + \left(\varepsilon + \frac{\log \ell}{n}\right) \\ \Lambda - \left(\varepsilon + \frac{\log \ell}{n}\right) \le \frac{1}{n} \log \frac{\|A^{\mp n}_h v\|}{\|v\|} &\le \Lambda + \left(\varepsilon + \frac{\log \ell}{n}\right) \end{aligned}$$

for all $v \in E^{\pm}_{h}$ and $n \ge 0$. The uniform convergence condition in the definition of uniformity guarantees that

$$\begin{aligned} -\Lambda - \varepsilon \le \frac{1}{n} \log \frac{\|A^{\pm n}_h v\|}{\|v\|} &\le -\Lambda + \varepsilon \\ \Lambda - \varepsilon \le \frac{1}{n} \log \frac{\|A^{\mp n}_h v\|}{\|v\|} &\le \Lambda + \varepsilon \end{aligned}$$

for large enough $n$, and we can always choose $\ell$ large enough to contain the transient behavior at small $n$.

To show that the angle condition is satisfied for large enough $\ell$, it's enough to prove that $d_{\angle}(E^{+}_{\alpha^n h}, E^{-}_{\alpha^n h})$ is bounded away from zero. The conditions in

Section 6.2 imply that $\ddot{H}$ is a compact metric space, the first return map on $\ddot{H}$ is minimal and uniquely ergodic, and the parallel transport cocycle $A$ is a uniform $\mathrm{SL}_2\,\mathbb{R}$ cocycle with nonzero Lyapunov exponent. Under these conditions, a classic result of Furman ensures that the stable distributions $E^{\pm}$ are continuous [7][9, Theorem 3]. Because $\ddot{H}$ is compact, and the distance between $E^+$ and $E^-$ can never be zero, it follows that this distance is bounded away from zero.

## 8.3 The stable distributions after abelianization

Recall that warping $\mathcal{E}$ along $\sigma$ has given us a new local system $\mathcal{F}$ and a stalkwise isomorphism $\Upsilon\colon \mathcal{E} \to \mathcal{F}$, supported on $\ddot{\Sigma}$. We're using the shorthand $F = \mathcal{F}_U$. Because $U$ is simple, the stalk restrictions of $\mathcal{E}$ and $\mathcal{F}$ identify $\mathcal{E}_p$ with $E$ and $\mathcal{F}_p$ with $F$ for every $p \in \ddot{U}$. We can thus view $\Upsilon$ as a map from $\ddot{U}$ to $\mathrm{SL}(E, F)$. Because $\sigma$ comes from a jump, $\Upsilon$ is constant along the vertical leaves of $\ddot{\Sigma}$, so in fact we can treat $\Upsilon$ as a map from $\ddot{H}$ to $\mathrm{SL}(E, F)$. This section takes place entirely within the flow box $U$, so we'll abbreviate $\sigma^U_{yx}$ as $\sigma_{yx}$ and $(y \mid x)^U$ as $(y \mid x)$.

Just as parallel transport in $\mathcal{E}$ along the vertical flow gave the linear cocycle $A\colon \ddot{H} \to \mathrm{SL}(E)$, parallel transport in $\mathcal{F}$ along the vertical flow gives a linear cocycle $\ddot{H} \to \mathrm{SL}(F)$. Define $F^{\pm}_h \subset F$ as the images of the lines $E^{\pm}_h$ under $\Upsilon_h$. Like $E^{\pm}$, the distributions $F^{\pm}$ are invariant under the parallel transport cocycle.[9] Let's put an inner product on $F$ by declaring $\Upsilon_a$, for some arbitrary $a \in \ddot{H}$, to be an isometry. We then get a sine metric on $\mathbf{P}F$, and we can ask whether the functions $F^{\pm}$ are Hölder.[10]

Recalling that $\Upsilon_h = \Upsilon_a \sigma_{ah}$, we see that $F^{\pm}_h = \Upsilon_a \sigma_{ah} E^{\pm}_h$ for all $h \in \ddot{H}$. Since $\Upsilon_a$ is, by definition, an isometry,

$$d_{\angle}(F^{\pm}_y, F^{\pm}_x) = d_{\angle}(\sigma_{ay} E^{\pm}_y, \sigma_{ax} E^{\pm}_x)$$

for all $y, x \in \ddot{H}$. We might therefore be able to prove that $F^{\pm}$ are Hölder by looking at how $\sigma^U$ affects distances in $\mathbf{P}E$.

## 8.4 The abelianized stable distributions are still Hölder

### 8.4.1 The deviation between nearby points is close to the identity

Remember the bound $C_{yx}$ we used in Section 7.3? We'll soon see that $\|\sigma_{yx} - 1\| \lesssim C_{yx}$. Combining this with the bound $C_{yx} \lesssim d(y, x)$ proven at the end of

[9] As a matter of fact, $\mathcal{F}$ should be uniform, with $F^{\pm}$ as its stable distributions. We don't need to know that, though.

[10] In light of the previous footnote, you might hope to show that $F^{\pm}$ are Hölder the same way we showed that $E^{\pm}$ are Hölder, by applying the theorem of Araújo, Bufetov, and Filip. The difficulty in this approach is that we don't know much about the parallel transport cocycle for $F$, making it hard to check the conditions of the theorem.

Section 7.3, we'll learn that

$$\|\sigma_{yx} - 1\| \lesssim d(y, x).$$

Let's get down to business. Recall that $\sigma_{yx}$ is the ordered product

$$\prod_{w \in (y|x)} s_w.$$

For any $w' \in W$, observe that

$$\begin{aligned}
\|\sigma_{yx} - 1\| &= \left\| \left( \prod_{w \in (y|x)} s_w \right) - 1 \right\| \\
&= \left\| \left( \prod_{w \in (y|x) \smallsetminus w'} s_w \right) + \left( \prod_{w \in (y|w')} s_w \right) (s_{w'} - 1) \left( \prod_{w \in (w'|x)} s_w \right) - 1 \right\| \\
&\leq \left\| \left( \prod_{w \in (y|x) \smallsetminus w'} s_w \right) - 1 \right\| + \|s_{w'} - 1\| \prod_{w \in (y|x) \smallsetminus w'} \|s_w\| \\
&\leq \left\| \left( \prod_{w \in (y|x) \smallsetminus w'} s_w \right) - 1 \right\| + \|s_{w'} - 1\| \prod_{w \in (y|x)} \|s_w\|.
\end{aligned}$$

Recalling the definition of $C_{yx}$, we see that $\prod_{w \in (y|x)} \|s_w\| \leq \exp C_{yx}$ by Proposition C.6.C. Thus,

$$\|\sigma_{yx} - 1\| \leq \left\| \left( \prod_{w \in (y|x) \smallsetminus w'} s_w \right) - 1 \right\| + \|s_{w'} - 1\| \exp C_{yx}.$$

By repeating the argument above, we see that for any finite subset $S$ of $(y \mid x)$,

$$\|\sigma_{yx} - 1\| \leq \left\| \left( \prod_{w \in (y|x) \smallsetminus S} s_w \right) - 1 \right\| + \left( \sum_{w \in S} \|s_w - 1\| \right) \exp C_{yx}.$$

The convergence of the product $\prod_{w \in (y|x)} s_w$ tells us that as $S$ grows, the first term of the inequality above goes to zero, leaving us with the bound

$$\begin{aligned}
\|\sigma_{yx} - 1\| &\leq \left( \sum_{w \in (y|x)} \|s_w - 1\| \right) \exp C_{yx} \\
&= C_{yx} \exp C_{yx}.
\end{aligned}$$

Since $C_{yx} \lesssim d(y,x)$, and distances in $\dot{\dot{H}}$ are bounded, we can bound $\exp C_{yx}$ by a constant. Hence, $\|\sigma_{yx} - 1\| \lesssim C_{yx}$. It follows, as explained at the beginning of the section, that

$$\|\sigma_{yx} - 1\| \lesssim d(y,x).$$

### 8.4.2 The abelianized stable distributions are Hölder

We know from Section 8.2 that $E^{\pm}$ are Hölder, say with exponent $\nu$. Because distances in $\dot{\dot{H}}$ are bounded, Hölder continuity with a given exponent implies Hölder continuity with all lower exponents, so we might as well assume for convenience that $\nu \leq 1$. For any $y, x \in \dot{\dot{H}}$, as discussed in Section 8.3,

$$\begin{aligned} d_{\angle}(F_y^{\pm}, F_x^{\pm}) &= d_{\angle}(\sigma_{ay} E_y^{\pm}, \sigma_{ax} E_x^{\pm}) \\ &= d_{\angle}(\sigma_{ay} E_y^{\pm}, \sigma_{ay}\sigma_{yx} E_x^{\pm}). \end{aligned}$$

Proposition D.D gives

$$d_{\angle}(F_y^{\pm}, F_x^{\pm}) \leq \|\sigma_{ya}\|^2 \, d_{\angle}(E_y^{\pm}, \sigma_{yx} E_x^{\pm}).$$

Because distances in $\dot{\dot{H}}$ are bounded, the result of the previous section ensures that $\|\sigma_{ya}\|$ is bounded as well. Therefore,

$$\begin{aligned} d_{\angle}(F_y^{\pm}, F_x^{\pm}) &\lesssim d_{\angle}(E_y^{\pm}, \sigma_{yx} E_x^{\pm}) \\ &\leq d_{\angle}(E_y^{\pm}, E_x^{\pm}) + d_{\angle}(E_x^{\pm}, \sigma_{yx} E_x^{\pm}). \end{aligned}$$

Proposition D.C combines with the bound from the previous section to show that

$$\begin{aligned} d_{\angle}(E_x^{\pm}, \sigma_{yx} E_x^{\pm}) &\leq \|\sigma_{yx} - 1\| \\ &\lesssim d(y,x). \end{aligned}$$

Using the fact that distances in $\dot{\dot{H}}$ are bounded and the assumption that $\nu \leq 1$, we conclude that

$$d_{\angle}(E_x^{\pm}, \sigma_{yx} E_x^{\pm}) \lesssim d(y,x)^{\nu}.$$

Meanwhile, the Hölder continuity of $E^{\pm}$ gives

$$d_{\angle}(E_y^{\pm}, E_x^{\pm}) \lesssim d(y,x)^{\nu}.$$

Therefore, altogether,

$$d_{\angle}(F_y^{\pm}, F_x^{\pm}) \lesssim d(y,x)^{\nu}.$$

In other words, the abelianized stable distributions $F^{\pm}$ are Hölder.

## 8.5 The abelianized stable distributions are constant

By construction, the values of the functions $F^{\pm}\colon \check{\check{H}} \to \mathbf{P}F$ match at adjacent edge points, as defined before Theorem 3.3.K. Since we just saw that $F^{\pm}$ are Hölder, and we're assuming the translation structure on $\Sigma$ satisfies the fat gap condition of Section 3.5.5, Corollary 3.3.L tells us that $F^{\pm}$ are constant.

Globally, this means the stable distributions $\mathcal{F}^{\pm}$ are constant with respect to the local system $\mathcal{F}$, so they're actually $\mathbb{R}^{\times}$ local subsystems of $\mathcal{F}$. In fact, $\mathcal{F}$ is the direct sum of the $\mathbb{R}^{\times}$ local systems $\mathcal{F}^{+}$ and $\mathcal{F}^{-}$.

# 9 A quick example

## 9.1 Overview

Now that we've proven abelianization works, let's see an example of what it does. The calculations in this section aren't rigorous, but I'll try to indicate what it would take to make them rigorous. Since abelianization is expected to work for $\mathrm{SL}_2\,\mathbb{C}$ cocycles, and in this case it does, we'll work over $\mathbb{C}$ rather than $\mathbb{R}$.

## 9.2 Setting the scene

### 9.2.1 A translation surface

Construct a torus with a translation structure by gluing the opposite sides of a parallelogram, inserting a singularity of cone angle $2\pi$ at the corner. For concreteness, let's fix the base of the parallelogram to be horizontal with length one, and set the height to be one as well. This leaves only one degree of freedom in the translation structure: the slope parameter $m$ labeled in the drawing below.

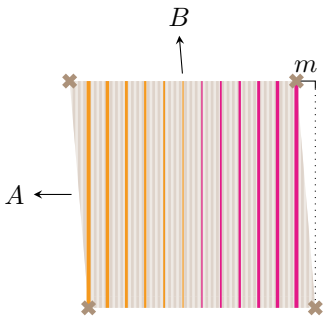


The torus has one forward-critical leaf and one backward-critical leaf. The drawing follows the critical leaves a little ways out from the singularity, so you can get an idea of how they wind around the surface. We'll assume $m$ is irrational, ensuring that neither critical leaf is a saddle connection.

The details of the computation depend on which way the parallelogram is leaning—a first hint of cluster-like behavior. To match the drawings, we'll show the work for the left-leaning case.

### 9.2.2 A variety of local systems

An $\mathrm{SL}_2\,\mathbb{C}$ local system on the torus minus the singularity is specified, up to isomorphism, by the group elements $A, B \in \mathrm{SL}_2\,\mathbb{C}$ that describe the parallel transport across the sides of the parallelogram, as shown in the drawing. Any values of $A$ and $B$ are possible.

Let's restrict ourselves to the dense open subset of the character variety in which $B$ has distinct eigenvalues. In this region, we can hit every isomorphism class using group elements of the form

$$A = \begin{bmatrix} \mu & \rho \\ \rho & \nu \end{bmatrix} \qquad\qquad B = \begin{bmatrix} \lambda & \cdot \\ \cdot & \frac{1}{\lambda} \end{bmatrix},$$

where $|\lambda| < 1$. Restricting further to the dense open subset in which $\mu\nu \notin (-\infty, 1]$, we can make $\rho$ a holomorphic function of $\mu$ and $\nu$ by noting that $\det A = 1$ and imposing the additional constraint $\mathrm{Re}\,\rho > 0$. This gives a holomorphic parameterization of a dense open subset of the character variety by the three variables $\mu$, $\nu$, and $\lambda$, which vary over the domain

$$\mu\nu \notin (-\infty, 1] \qquad\qquad |\lambda| < 1.$$

### 9.2.3 Reduction to an interval exchange

Let $Z$ be the horizontal segment running across the middle of the parallelogram. Under the vertical flow, $Z$ sweeps out a simple flow box that covers almost the whole torus. The vertical edge of the flow box is non-critical, so we can compute the abelianized local system just by looking at the parallel transport cocycle over $Z$. For this purpose, we'll mostly carry on with the notation from Section 7.

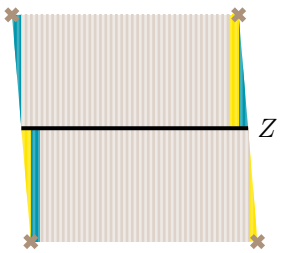


Identify $Z$ with $(-1, 0)$. The first return relation $\alpha$ has a single break point, $b = -\frac{m}{2}$. Its inverse $\alpha^{-1}$ has break point $c = -1 + \frac{m}{2}$. Because $Z$ isn't well-cut, the break points aren't the only points where $\alpha$ and $\alpha^{-1}$ return nothing: $-m$ and $-1 + m$ vanish under the actions of $\alpha$ and $\alpha^{-1}$ as well. We aren't calling the latter break points because they don't lie on critical leaves.

The forward parallel transport cocycle is constant on the intervals

$$\left(-1,-m\right) \qquad \left(-m,-\tfrac{m}{2}\right) \qquad \left(-\tfrac{m}{2},0\right),$$

where it has the values

$$B \qquad A^{-1}B \qquad BA^{-1},$$

respectively.

## 9.3 Abelianization

### 9.3.1 Approximation

There isn't an obvious way to compute the abelianized local system exactly, but there is a pretty obvious way to approximate it when $m$ is tiny. As before, we'll show the work for the left-leaning case.

### 9.3.2 The slithering jumps at the break points

As $m$ approaches zero, the sequence of $\mathrm{SL}_2\,\mathbb{C}$ elements generated by repeatedly applying the forward parallel transport cocycle to $\overrightarrow{b}$ approaches

$$\ldots B, B, B, B, BA^{-1},$$

in the sense that it takes more and more iterations to deviate from this sequence. Using the discussion after Lemma 2.2 of [8], you can deduce from this that the forward-stable line $E^+_{\overrightarrow{b}}$ approaches the line spanned by

$$A\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} \mu \\ \rho \end{bmatrix} \sim \begin{bmatrix} 1 \\ \frac{\rho}{\mu} \end{bmatrix}$$

Similarly, applying the forward parallel transport cocycle to $\overleftarrow{b}$ yields a sequence approaching

$$\ldots B, B, B, B, A^{-1}B$$

as $m$ goes to zero, so $E^+_{\overleftarrow{b}}$ approaches the span of

$$B^{-1}A\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{\lambda}\,\mu \\ \lambda\,\rho \end{bmatrix} \sim \begin{bmatrix} 1 \\ \lambda^2\,\frac{\rho}{\mu} \end{bmatrix}.$$

On the other hand, applying the backward parallel transport cocycle to $b$ yields a sequence approaching

$$\ldots B^{-1}, B^{-1}, B^{-1}, B^{-1}, B^{-1},$$

so $E^-_b$ goes to the span of

$$\begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

The slithering jump $s_b$ therefore approaches

$$\begin{bmatrix} & 1 & \cdot \\ (\lambda^2-1)\frac{\rho}{\mu} & & 1 \end{bmatrix}$$

as $m$ goes to zero. A similar computation shows that $s_c$ approaches

$$\begin{bmatrix} 1 & (\lambda^2-1)\frac{\rho}{\mu} \\ \cdot & & 1 \end{bmatrix}$$

in the same limit.

### 9.3.3 The slithering jumps at all the critical points

The slithering jumps at all the critical points can be deduced from the ones at the break points using the flow-invariance property discussed in Section 6.3.4. For each $n \geq 0$, the jump $s_{\alpha^{-n}b}$ at the forward-critical point $\alpha^{-n}b$ goes to

$$B^{-n}s_bB^n = \begin{bmatrix} & 1 & \cdot \\ \lambda^{2n}(\lambda^2-1)\frac{\rho}{\mu} & & 1 \end{bmatrix}$$

as $m$ goes to zero. The jump $s_{\alpha^n c}$ at the backward-critical point $\alpha^n c$ goes to

$$B^n s_c B^{-n} = \begin{bmatrix} 1 & \lambda^{2n}(\lambda^2-1)\frac{\rho}{\mu} \\ \cdot & & 1 \end{bmatrix}$$

in the same limit.

### 9.3.4 The slithering deviation

Looking back at the drawing in Section 9.2.1, you can see that forward-critical points $\alpha^{-n}b$ march from right to left across $Z$ as $n$ grows, while the backward-critical points $\alpha^n c$ march from left to right. As $m$ goes to zero, it takes longer and longer for the parades to meet at $-\frac{1}{2} \in Z$. It seems clear that the bound from Section 7.2 will hold more or less uniformly as $m$ goes to zero, so we shouldn't have to worry too much about the later jumps. We can therefore compute as though the parades never meet.

In this approximation, the deviation to $-\frac{1}{2}$ from 0 is given by the product

$$\cdots \quad s_{\alpha^{-3}b} \quad s_{\alpha^{-2}b} \quad s_{\alpha^{-1}b} \quad s_b,$$

and the deviation to $-1$ from $-\frac{1}{2}$ is given by

$$s_c \quad s_{\alpha^1 c} \quad s_{\alpha^2 c} \quad s_{\alpha^3 c} \quad \cdots .$$

As $m$ goes to zero, the jumps $s_{\alpha^{-n}b}$ go to commuting shears, so you should be able to show that their product approaches

$$\begin{bmatrix} 1 & & \cdot \\ \frac{\rho}{\mu}(\lambda^2-1)\sum_{n=0}^{\infty}\lambda^{2n} & & 1 \end{bmatrix} = \begin{bmatrix} 1 & \cdot \\ -\frac{\rho}{\mu} & 1 \end{bmatrix}$$

Similarly, the product of the jumps $s_{\alpha^n c}$ should approach

$$\begin{bmatrix} 1 & \frac{\rho}{\mu}(\lambda^2-1)\sum_{n=0}^{\infty}\lambda^{2n} \\ \cdot & 1 \end{bmatrix} = \begin{bmatrix} 1 & -\frac{\rho}{\mu} \\ \cdot & 1 \end{bmatrix}$$

Now we know enough to approximate the holonomy $A_{\text{ab}}$ of the abelianized local system around the loop that starts at $-\frac{1}{2}$, runs left to 0, wraps around to 1, and runs left back to $-\frac{1}{2}$. As $m$ goes to zero, this holonomy approaches

$$\begin{bmatrix} 1 & \cdot \\ -\frac{\rho}{\mu} & 1 \end{bmatrix} A \begin{bmatrix} 1 & -\frac{\rho}{\mu} \\ \cdot & 1 \end{bmatrix} = \begin{bmatrix} \mu & \cdot \\ \cdot & \frac{1}{\mu} \end{bmatrix}.$$

As expected, the abelianized holonomy preserves the stable lines $E^+_{-1/2}$ and $E^-_{-1/2}$, which approach

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

as $m$ goes to zero. In the limit, abelianization has no effect on the holonomy around a vertical loop, so $B_{\text{ab}}$ goes to $B$ as $m$ goes to zero.

## 9.4 Spectral coordinates

We've learned that on the translation torus constructed from a left-leaning parallelogram with slope parameter $m$, the abelianization of the local system

$$A = \begin{bmatrix} \mu & \rho \\ \rho & \nu \end{bmatrix} \qquad\qquad B = \begin{bmatrix} \lambda & \cdot \\ \cdot & \frac{1}{\lambda} \end{bmatrix}$$

approaches

$$A_{\text{ab}} = \begin{bmatrix} \mu & \cdot \\ \cdot & \frac{1}{\mu} \end{bmatrix} \qquad\qquad B_{\text{ab}} = \begin{bmatrix} \lambda & \cdot \\ \cdot & \frac{1}{\lambda} \end{bmatrix}$$

as $m$ goes to zero. For a right-leaning parallelogram, the analogous calculation shows that the abelianization approaches

$$A_{\text{ab}} = \begin{bmatrix} \frac{1}{\nu} & \cdot \\ \cdot & \nu \end{bmatrix} \qquad\qquad B_{\text{ab}} = \begin{bmatrix} \lambda & \cdot \\ \cdot & \frac{1}{\lambda} \end{bmatrix}$$

as $m$ goes to zero.

As expected, the abelianized local system splits into a pair of $\mathbb{C}^\times$ local systems comprising the forward- and backward-stable lines of the original. Restricing our attention to the forward-stable local systems, we get the limiting holonomies

$$A^+_{\text{ab}} = \mu \qquad\qquad B^+_{\text{ab}} = \lambda$$

in the left-leaning case, and

$$A^+_{\text{ab}} = \tfrac{1}{\nu} \qquad\qquad B^+_{\text{ab}} = \lambda$$

in the right-leaning case. Looking at both limits, we can recover the holomorphic coordinates $\mu$, $\nu$, and $\lambda$ that we've been using to parameterize a dense open subset of the $\mathrm{SL}_2\,\mathbb{C}$ character variety.

# 10 Future directions

## 10.1 Abelianization should be more

The main result of this paper has been to show that, for a generic compact translation surface $\Sigma$, abelianization gives a well-defined map from an open subspace of the $\mathrm{SL}_2\,\mathbb{R}$ local systems on $\Sigma \smallsetminus \mathfrak{B}$ to the space of $\mathbb{R}^\times$ local systems on $\Sigma \smallsetminus \mathfrak{B}$. In the previously studied case of abelianization on a punctured half-translation surface, however, the mere existence of the abelianization map is the least interesting of its properties. Many of the special features of the abelianization map are expected to persist for compact half-translation surfaces.

Let's review the setting for this version of abelianization. Since we're speculating anyway, let's assume the results of this paper can be extended to $\mathrm{SL}_2\,\mathbb{C}$ local systems, as discussed in Section 1.2.3. Suppose $\Sigma$ is the translation double cover of a compact half-translation surface $C$. Recall from Section 1.1 the definition of the twisted $\mathrm{SL}_2\,\mathbb{C}$ character variety $\mathcal{M}_{-1}C$. Pushback of local systems along the covering $\Sigma \to C$ gives a map from $\mathcal{M}_{-1}C$ to $\mathcal{M}_1\Sigma$. Composing this map with the abelianization process, we get an abelianization map that sends twisted $\mathrm{SL}_2\,\mathbb{C}$ local systems on $C$ to $\mathbb{C}^\times$ local systems on $\Sigma \smallsetminus \mathfrak{B}$. The abelianized local systems turn out to have holonomy $-1$ around each singularity, so abelianization actually sends twisted $\mathrm{SL}_2\,\mathbb{C}$ local systems on $C$ to twisted $\mathbb{C}^\times$ local systems on $\Sigma$.

## 10.2 It should be a Darboux chart on $\mathcal{M}_{-1}C$

When the compact translation surface $C$ is replaced with a punctured half-translation surface $C'$, Gaiotto, Moore, and Neitzke showed that the abelianization map is a symplectomorphism onto its image [5, §§ 10.4 and 10.8]. It can therefore be seen as a Darboux coordinate system on the dense open subset of $\mathcal{M}_{-1}C'$ where it's defined. This is the *spectral coordinate system* mentioned in Section 1.1. In the compact case, we've only managed to define abelianization for uniform local systems, so its domain is no longer expected to be dense. The other properties, however, should persist, and they've already been proven in the parallel work of Bonahon and Dreyer: [6, §6] shows that the map sending a Hitchin character to its shearing cycle is invertible, and the discussion in [6, §7] should imply that this map is a symplectomorphism with respect to some natural symplectic structure on the space of tangent cycles. It may be interesting to see how the analogous arguments play out for our construction, both for the sake of learning how the objects considered by Bonahon and Dreyer appear in our setting, and also in case the domain of the abelianization map turns out not to be precisely analogous to the domain of the shearing cycle.

## 10.3 It may be a generalized cluster coordinate chart

In the punctured case, Gaiotto, Moore, and Neitzke showed that the spectral coordinates are actually Fock-Goncharov coordinates. In particular, the spectral

coordinate systems produced by rotating the half-translation structure of $C'$ fit together into a cluster algebra. In the compact case, rotationally related spectral coordinate systems also appear to fit together into something bigger, but that something doesn't seem to be a cluster algebra: it has mutation-like behavior, but no readily identifiable clusters. Some preliminary investigations of this structure are briefly reported in [23].

## 10.4 It should be holomorphic

On a punctured half-translation surface, the source and target of the abelianization map are not only symplectic manifolds, but holomorphic symplectic manifolds, and abelianization is a holomorphic symplectomorphism [5, §10.4]. The holomorphicity of the abelianization map isn't special to half-translation surfaces, and it appears to persist in the compact case. Here's a sketch of an argument that the abelianization map for a compact translation surface $\Sigma$ is holomorphic. Note that $\Sigma$ doesn't have to be the translation double cover of a half-translation surface.

An easy way to see the complex structure on the moduli space of $\mathrm{SL}_2\,\mathbb{C}$ local systems on $\Sigma \smallsetminus \mathfrak{B}$ is to pick a well-cut flow box in $\overset{\leftrightarrow}{\Sigma}$. As we observed in Section 5.5, a local system on $\Sigma \smallsetminus \mathfrak{B}$ is described up to isomorphism by an interval cocycle over the flow box's first return relation $\alpha$, and interval cocycles cut across all the isomorphism classes. An interval cocycle over $\alpha$ is just an element $A$ of $(\mathrm{SL}_2\,\mathbb{C})^{\mathcal{A}^+}$, where $\mathcal{A}^+$ is the set of intervals exchanged by $\alpha$. The complex structure of this space matches the one on the space of local systems.

Suppose $A$ is uniform. As in Section 6.3, write the slithering jump for $A$ at a critical point $w$ of $\alpha$ as $s_w \in \mathrm{SL}_2\,\mathbb{C}$. The deviation $\sigma$ that abelianizes $A$ is given by ordered products of $s_w$ over the critical points. The Lyapunov exponent of $A$ can be bounded when $A$ varies over a small enough region [8, proof of Proposition 2.6], so the bound on the size of $s_w$ in Section 7.2 should be uniform with respect to $A$. The products that define $\sigma$ should therefore converge uniformly with respect to $A$. If this works, the task of showing that $\sigma$ varies holomorphically with $A$ is reduced to the task of showing that $s_w$ does for each critical point $w$.

Say $w$ is a backward-critical point. The formulas in Section 6.3 make it clear that $s_w$ depends holomorphically on the stable lines $E^-_{\overleftarrow{w}}, E^+_w, E^-_{\overrightarrow{w}} \in \mathbf{P}\mathbb{C}^2$. The line $E^\pm_s$ is approximated by the line most contracted by $A^{\pm n}_s$ when $n$ is large [8, proof of Proposition 2.1]. In fact, the convergence of the most contracted line to the stable line is uniform with respect to $A$ [8, discussion around Equation 2.6], so we just need to show that the most contracted line depends holomorphically on $A$.

Using the standard inner product on $\mathbb{C}^2$, we can observe that the line most contracted by $A^{\pm n}_s$ lies in the eigenspace of $(A^{\pm n}_s)^\dagger(A^{\pm n}_s)$ with the smallest eigenvalue. Since $A$ is uniform, it should be safe to assume that $(A^{\pm n}_s)^\dagger(A^{\pm n}_s)$ has distinct eigenvalues for large enough $n$, so the most contracted line is just the eigenline with the smallest eigenvalue. The relationship between an operator

and its eigenlines is holomorphic, and the map that sends $A$ to $A_s^{\pm n}$ is too, so we should be done.

# A Technical tools for warping local systems

## A.1 The lily pad lemma

**Lemma A.1.A.** *Suppose $\mathcal{U}$ is an open cover of a connected space. For any two points $a$ and $b$ in the space, there is a finite sequence of elements of $\mathcal{U}$ in which the first element contains $a$, the last element contains $b$, and every element intersects the next one.*

*Proof.* Let's call a finite sequence of elements of $\mathcal{U}$ a *lily path* if every element intersects the next one. We'll say two points $a$ and $b$ can be "connected by a lily path" if there's a lily path whose first element contains $a$ and whose last element contains $b$.

A lily path connecting $a$ to $b$ also connects $a$ to every other point in the last element of the path, which is an open neighborhood of $b$. Hence, the set of points that can be connected to $a$ by a lily path is open.

On the other hand, suppose $b$ can't be connected to $a$ by a lily path. Since $\mathcal{U}$ is a cover, there's some $U \in \mathcal{U}$ containing $b$. If a point in $U$ could be connected to $a$ by a lily path, adding $U$ to the end of that path would give a lily path connecting $a$ to $b$. Hence, no point in $U$ can be connected to $a$ by a lily path. Therefore, the set of points that can't be connected to $a$ by a lily path is also open. □

## A.2 Collapsing downward-directed colimits

**Proposition A.2.A.** *Suppose $\Lambda$ is a downward-directed set, $C$ is some category, and $F\colon \Lambda \to C$ is a diagram in which all the arrows are isomorphisms. Then $F$ has a colimit, and the defining arrows from the diagram to its colimit are isomorphisms.*

*Proof.* Pick any object $s$ of $\Lambda$. For any other object $t$, we can get an isomorphism $f_t\colon F(t) \to F(s)$ by picking a common lower bound $\check{t}$ of $s$ and $t$ and taking the composition $F(\check{t} \leftarrow s)^{-1} F(\check{t} \leftarrow t)$. The isomorphism we get doesn't depend on our choice of $\check{t}$.

Now, pick any object $c$ of $C$. Suppose that for each object $t$ of $\Lambda$, we have an arrow $\phi_t\colon F(t) \to c$, and these arrows commute with the arrows of the diagram $F$. Observing that

$$\begin{aligned}\phi_s f_t &= \phi_s F(\check{t} \leftarrow s)^{-1} F(\check{t} \leftarrow t) \\ &= \phi_{\check{t}} F(\check{t} \leftarrow t) \\ &= \phi_t,\end{aligned}$$

we see that the object $F(s)$, equipped with the isomorphisms $f_t\colon F(t) \to F(s)$, is a colimit of the diagram $F$. □

# B Relational dynamics

Some dynamical systems, including vertical flows on singular translation surfaces, interval exchanges, and even the humble doubling map, are discontinuous if you insist on defining them at every point. (These particular examples are discussed in Sections 3.2.1 and 3.3.2.) You can recover a kind of continuity, however, if you describe the dynamics using relations rather than maps, allowing points to get lost as they fall into singularities, breaks between intervals, or whatever.

Consider a relation $\phi$ between topological spaces $Y$ and $X$. Define

$$\phi A = \{y \in Y : y \,\phi\, a \text{ for some } a \in A\}$$
$$B\phi = \{x \in X : b \,\phi\, x \text{ for some } b \in B\}$$

for subsets $A \subset X$ and $B \subset Y$. For convenience, we'll relax the distinction between singletons and points, denoting $\phi\{x\}$, for example, by $\phi x$. If $\phi$ is a function $Y \leftarrow X$, then $\phi x \in Y$ is the value of $\phi$ at a point $x \in X$, and $B\phi \subset X$ is the preimage of a subset $B \subset Y$.

Define $\phi$ to be *injective* if $y\phi$ is a singleton for all $y \in Y$, *coinjective* if $\phi x$ is a singleton for all $x \in X$, and *biinjective* if it's both injective and coinjective. In less baroque language, a coinjective relation is just a partially defined function.

Define $\phi$ to be *continuous* if $V\phi$ is open whenever $V \subset Y$ is open, *cocontinuous* if $\phi U$ is open whenever $U \subset X$ is open, and *bicontinuous* if it's both continuous and cocontinuous. If $\phi$ is a function, "continuous" means what it usually means, and "cocontinuous" means "open." Local systems can be pushed forward and backward along a bicontinuous relation.

A *flow by bicontinuous relations* on $X$ can be defined as a relation $\psi$ between $X$ and $\mathbb{R} \times X$ with the following properties:

- As a whole, $\psi$ is continuous and coinjective.
- For all $t \in \mathbb{R}$, the relation $\psi^t = \,_\, \psi\, (t, _)$ is bicontinuous and biinjective.
- For all $t, s \in \mathbb{R}$, we have $\psi^t \psi^s = \psi^{t+s}$.

This kind of partially defined flow acts a lot like an ordinary flow by homeomorphisms. In particular, if $X$ carries a local system $\mathcal{E}$, it gives a parallel transport morphism $\mathcal{E}_{\psi^t U} \leftarrow \mathcal{E}_U$ for every open subset $U$ of $X$ and every time $t$.

# C Infinite ordered products

## C.1 Definition

Suppose $M$ is a Hausdorff topological monoid, like the one formed by the endomorphisms or automorphisms of a finite-dimensional vector space, and $A$ is

a totally ordered set. Given a function $m\colon A \to M$, we'd like to make sense of the potentially infinite ordered product

$$\prod_{p \in A} m_p,$$

which I'll refer to in writing as "the product of $m$ over $A$."

Recall that a function from a directed set $\Lambda$ into a topological space is called a *net*. For any $s \in \Lambda$, let's call the set $\{t \in \Lambda : t \geq s\}$ the *shadow* of $s$. A net is said to *converge* to a point if, for every open neighborhood $\Omega$ of that point, there is some element of $\Lambda$ whose shadow is sent by the net into $\Omega$. A net into a Hausdorff space, like $M$, converges to at most one point.

The finite subsets of $A$ form a directed set under inclusion, and the product of $m$ over any finite subset is well-defined, so we can define the product of $m$ over $A$ to be the limit of the net that sends each finite subset of $A$ to the product of $m$ over that subset. I'll call this net the "product net" for short.

## C.2 Calculation

Suppose $\Lambda'$ and $\Lambda$ are directed sets, $n$ is a net on $\Lambda$, and $f\colon \Lambda' \to \Lambda$ is an order-preserving map whose image intersects the shadow of every element of $\Lambda$. In these circumstances, the net $n \circ f$ is called a *subnet* of $n$. If $n$ converges to a certain point, every subnet of $n$ converges to that point.

In particular, let $A$ be a totally ordered set, and $\Lambda$ its finite subsets. If $f\colon \mathbb{N} \to \Lambda$ is an increasing sequence of subsets whose union is all of $A$, the image of $f$ intersects the shadow of every finite subset. Thus, if a product over $A$ converges, we can find it by looking at partial products over any sequence of finite subsets whose union is $A$. Such a sequence must exist if $A$ is countable.

## C.3 Composition

Given two totally ordered sets $B$ and $A$, let $B \sqcup A$ be the disjoint union of $B$ and $A$ with the total order that makes the inclusions order-preserving and puts every element of $B$ to the left of every element of $A$.

**Proposition C.3.A.** *Say we have a function $m\colon B \sqcup A \to M$. If the products of $m$ over $B$ and $A$ converge, then the product over $B \sqcup A$ converges too, and*

$$\prod_{p \in B \sqcup A} m_p = \left(\prod_{p \in B} m_p\right)\left(\prod_{p \in A} m_p\right).$$

*Proof.* Let $\beta$ and $\alpha$ be the products of $m$ over $B$ and $A$, respectively. Given an open neighborhood $\Omega$ of $\beta\alpha$, we want to find a finite subset of $B \sqcup A$ whose shadow is sent by the product net into $\Omega$.

Using the fact that multiplication in $M$ is continuous, find open neighborhoods $\Omega_B$ of $\beta$ and $\Omega_A$ of $\alpha$ with $\Omega_B\Omega_A \subset \Omega$. Then, find finite subsets $S_B$ and $S_A$ of $B$ and $A$ whose shadows are sent into $\Omega_B$ and $\Omega_A$, respectively.

For any finite subset $R$ of $B \sqcup A$ containing $S_B \cup S_A$, we can use the fact that $B > A$ to rewrite the product of $m$ over $R$ as

$$\left(\prod_{p \in R \cap B} m_p\right)\left(\prod_{p \in R \cap A} m_p\right).$$

Observing that $R \cap B$ contains $S_B$ and $R \cap A$ contains $S_A$, we conclude that the product of $m$ over $R$ is in $\Omega$. □

## C.4 Equivariance

**Proposition C.4.A.** *If the product of $m\colon A \to M$ over $A$ converges,*

$$\phi\left(\prod_{p \in A} m_p\right) = \prod_{p \in A} \phi(m_p)$$

*for any continuous homomorphism $\phi\colon M \to M$.*

*Proof.* Given an open neighborhood $\Omega$ of the left-hand side, we want to find a finite subset $S$ of $A$ with the property that

$$\prod_{p \in R} \phi(m_p) \in \Omega$$

for all finite subsets $R \subset A$ containing $S$.

Since $\phi$ is continuous, $\phi^{-1}(\Omega)$ is open neighborhood of the product of $m$ over $A$. By the definition of convergence, we can find a finite subset $S$ of $A$ with the property that

$$\prod_{p \in R} m_p \in \phi^{-1}(\Omega)$$

for all finite subsets $R \subset A$ containing $S$. Applying $\phi$ to both sides, we see that $S$ is just what we wanted. □

## C.5 Inversion

Instead of just a map into a topological monoid, suppose we have a map $g\colon A \to G$ into a topological group. Let $A^{\mathrm{op}}$ be $A$ with the opposite order.

**Proposition C.5.A.** *If the product of $g$ over $A$ converges,*

$$\prod_{p \in A^{\mathrm{op}}} g_p^{-1} = \left(\prod_{p \in A} g_p\right)^{-1}.$$

*Proof.* Analogous to the proof of Proposition C.4.A. □

We won't use this result for anything. It's only here to reassure you that the directionality of our construction of deviations from jumps in Section 4.3 doesn't introduce any actual asymmetry.

## C.6 Convergence

Any Banach algebra, like $\operatorname{End}\mathbb{R}^2$ with the operator norm, can be thought of as a topological monoid by forgetting the addition. Since all the ordered products we care about will be taken in $\mathrm{SL}_2\,\mathbb{R}$, a closed submonoid of $\operatorname{End}\mathbb{R}^2$, understanding ordered products in a Banach algebra will be very helpful to us. The following results generalize Theorem 2.3, Corollary 2.4, and a simplified version of Theorem 2.7 from [24] to products over arbitrarily ordered index sets.

**Proposition C.6.A.** *Suppose we have a totally ordered set $A$, a unital Banach algebra $\mathcal{X}$, and a function $x\colon A \to \mathcal{X}$. If the sum $\sum_{p\in A} \|x_p - 1\|$ converges, the product $\prod_{p\in A} x_p$ converges as well.*

**Proposition C.6.B.** *If the sum in the proposition above converges, and each factor $x_p$ is invertible, the product is invertible.*

Our proof will give a handy bound for free.

**Proposition C.6.C.** *If the sum in Proposition C.6.A converges, the product $\prod_{p\in A} \|x_p\|$ converges as well, and*

$$\left\|\prod_{p\in A} x_p\right\| \quad\leq\quad \prod_{p\in A} \|x_p\| \quad\leq\quad \exp\left(\sum_{p\in A} \|x_p - 1\|\right).$$

Let's start with a less ambitious result.

**Proposition C.6.D.** *If the sum in Proposition C.6.A converges, then for any $C > \exp\left(\sum_{p\in A} \|x_p - 1\|\right)$, there's a finite subset $S$ of $A$ with the property that*

$$\prod_{p\in R} \|x_p\| < C$$

*for all finite subsets $R \subset A$ containing $S$.*

*Proof.* Assume $\sum_{p\in A} \|x_p - 1\|$ converges, and consider any positive constant $C$ with $\log C > \sum_{p\in A} \|x_p - 1\|$. By the definition of convergence, we can find a finite subset $S$ of $A$ with the property that

$$\sum_{p\in R} \|x_p - 1\| < \log C$$

for all finite subsets $R \subset A$ containing $S$. Since

$$\log \|x_p\| \leq \big|\|x_p\| - 1\big| \qquad \text{and} \qquad \big|\|x_p\| - 1\big| \leq \|x_p - 1\|$$

for all $p \in A$, we know

$$\sum_{p\in R} \log \|x_p\| < \log C.$$

for all finite subsets $R \subset A$ containing $S$. Exponentiating both sides gives the desired result. □

*Proof of Proposition C.6.A.* By definition, $\mathcal{X}$ is complete, so we can prove that the product converges by showing that the product net is Cauchy. To that end, given any $\epsilon > 0$, we want to find a finite subset $S$ of $A$ with the property that

$$\left\| \prod_{p\in R} x_p - \prod_{p\in S} x_p \right\| < \epsilon$$

for all finite subsets $R \subset A$ containing $S$.

Assume the sum $\sum_{p\in A} \|x_p - 1\|$ converges, and pick a constant $C > \exp\left(\sum_{p\in A} \|x_p - 1\|\right)$. Every convergent net is Cauchy [25, Proposition 3.2], so we can find a finite subset $S'$ of $A$ with the property that

$$\left| \sum_{p\in R} \|x_p - 1\| - \sum_{p\in S'} \|x_p - 1\| \right| < \epsilon/C$$

for all finite subsets $R \subset A$ containing $S'$. This inequality simplifies to

$$\sum_{p\in R\smallsetminus S'} \|x_p - 1\| < \epsilon/C.$$

By Proposition C.6.D, we can also find a finite subset $S''$ of $A$ with the property that

$$\prod_{p\in R} \|x_p\| < C$$

for all finite subsets $R \subset A$ containing $S''$. Defining $S$ as $S' \cup S''$, and observing that $R \smallsetminus S'$ contains $R \smallsetminus S$, we see that

$$\sum_{p\in R\smallsetminus S} \|x_p - 1\| < \epsilon/C \qquad \text{and} \qquad \prod_{p\in R} \|x_p\| < C$$

for all finite subsets $R \subset A$ containing $S$.

Put the elements of $R \smallsetminus S$ in some order $r_1, \ldots, r_n$. Let

$$\begin{aligned} R_0 &= S, \\ R_1 &= R_0 \cup \{r_1\}, \\ R_2 &= R_1 \cup \{r_2\}, \end{aligned}$$

and so on. Let

$$\Delta_{k+1} = \left\| \prod_{p\in R_{k+1}} x_p - \prod_{p\in R_k} x_p \right\|$$

Notice that

$$\prod_{p\in R_{k+1}} x_p - \prod_{p\in R_k} x_p = \left( \prod_{\substack{p\in R_k \\ p>r_{k+1}}} x_p \right) (x_{r_{k+1}} - 1) \left( \prod_{\substack{p\in R_k \\ r_{k+1}>p}} x_p \right),$$

yielding the bound

$$\begin{aligned}\Delta_{k+1} &\leq \|x_{r_{k+1}} - 1\| \prod_{p \in R_k} \|x_p\| \\ &\leq \|x_{r_{k+1}} - 1\| C.\end{aligned}$$

Finally, observe that

$$\begin{aligned}\left\| \prod_{p \in R} x_p - \prod_{p \in S} x_p \right\| &\leq \Delta_1 + \ldots + \Delta_n \\ &\leq \|x_{r_1} - 1\| C + \ldots + \|x_{r_n} - 1\| C \\ &= C \sum_{p \in R \smallsetminus S} \|x_p - 1\| \\ &< \epsilon.\end{aligned}$$

Since $R$ could have been any finite subset of $A$ containing $S$, and the method we used to find $S$ works for any $\epsilon > 0$, we've proven that the product net is Cauchy. □

*Proof of Proposition C.6.B.* Given an ordered set $I$, let $I^{\mathrm{op}}$ be the same set in the opposite order. Assume $\sum_{p \in A} \|x_p - 1\|$ converges. Equivalently, because addition is commutative, $\sum_{p \in A^{\mathrm{op}}} \|x_p - 1\|$ converges. This is only possible if $\|x_p - 1\| \leq 6/7$ for all but finitely many $p \in A^{\mathrm{op}}$. It follows, by the calculation in the proof of [24, Lemma 2.6], that $\|x_p^{-1}\| \leq 7$ for all but finitely many $p \in A^{\mathrm{op}}$. That means $\|x_p^{-1}\|$ has a maximum over all $p \in A^{\mathrm{op}}$, which I'll call $M$. Just as in the proof of [24, Theorem 2.7], observe that

$$\begin{aligned}\|x_p^{-1} - 1\| &= \|x_p^{-1}(1 - x_p)\| \\ &\leq M \|x_p - 1\|\end{aligned}$$

for all $p \in A^{\mathrm{op}}$. Then [26, Exercise 7.40.c] tells us that $\sum_{p \in A^{\mathrm{op}}} \|x_p^{-1} - 1\|$ converges, so $\prod_{p \in A^{\mathrm{op}}} x_p^{-1}$ converges by Proposition C.6.A.

Define $y'' = \prod_{p \in A^{\mathrm{op}}} x_p^{-1}$ and $y' = \prod_{p \in A} x_p$. Multiplication in a Banach algebra is continuous, so for any open neighborhood $\Omega$ of $y''y'$, we can find open neighborhoods $\Omega''$ of $y''$ and $\Omega'$ of $y'$ with $\Omega''\Omega' \subset \Omega$. By convergence, we can find finite subsets $S''$ and $S'$ of $A$ such that

$$\prod_{p \in R^{\mathrm{op}}} x_p^{-1} \in \Omega''$$

for all finite subsets $R \subset A$ containing $S''$, and

$$\prod_{p \in R} x_p \in \Omega'$$

for all finite subsets $R \subset A$ containing $S'$. Now, defining $S$ as $S'' \cup S'$, we can observe that

$$\left(\prod_{p \in S^{\mathrm{op}}} x_p^{-1}\right)\left(\prod_{p \in S} x_p\right) \in \Omega.$$

But the product in the expression above is clearly equal to 1! We've shown that every open neighborhood of $y''y'$ contains 1, which means $y''y'$ is equal to 1. The same argument can be used to show that $y'y''$ is 1. Therefore, $\prod_{p \in A} x_p$ is invertible, with inverse $\prod_{p \in A^{\mathrm{op}}} x_p^{-1}$. □

*Proof of Proposition C.6.C.* Assume $\sum_{p \in A} \|x_p - 1\|$ converges. By Proposition C.6.A, $\prod_{p \in A} x_p$ converges too. The convergence of $\prod_{p \in A} \|x_p\|$ follows immediately from the fact that the norm is continuous. The first inequality is easy to establish.

Because $\sum_{p \in A} \|x_p - 1\|$ is a sum of non-negative numbers, its convergence implies that at most countably many of its terms are nonzero. We can therefore assume, without loss of generality, that $A$ is countable. Combining this fact with Proposition C.6.D and the discussion in Section C.2, it's not hard to show that $\prod_{p \in A} \|x_p\|$ is less than any number greater than $\exp\left(\sum_{p \in A} \|x_p - 1\|\right)$. That gives the second inequality. □

# D Linear algebra on the Euclidean plane

The Euclidean plane is a two-dimensional real inner product space $E$ with a volume form $D$ in which the unit square has unit volume. It will be useful to collect some basic facts about geometry in such a space. Really well-known facts will be stated without proof.

The sine metric $d_\angle$ on $\mathbf{P}E$ is defined as in Section 8.2. The area of a parallelogram can be computed from the lengths of its sides and the angle between them.

**Proposition D.A.**

$$|D(u, v)| = \|u\|\|v\| \, d_\angle(u, v)$$

*for all $u, v \in E$.*

By comparing the areas of some well-chosen parallelograms, you can deduce the law of sines.

**Proposition D.B** (The law of sines)**.**

$$\|u\| \, d_\angle(u, u + v) = \|v\| \, d_\angle(v, u + v)$$

*for all $u, v \in E$ with $u + v \neq 0$.*

From this we see that the extent to which a linear map $T$ can move lines is limited by the operator norm of $T - 1$.

**Proposition D.C.** *For any linear map $T\colon E \to E$,*

$$d_\angle(u, Tu) \le \|T - 1\|$$

*whenever $Tu \neq 0$.*

*Proof.* Because distances in $\mathbf{P}F$ are never greater than one, it follows from the law of sines that

$$d_\angle(u, u+v) \le \frac{\|v\|}{\|u\|}$$

In particular,

$$d_\angle(u, Tu) \le \frac{\|(T-1)u\|}{\|u\|}.$$

□

The extent to which a volume-preserving map can expand angles is limited by the operator norm of its inverse.

**Proposition D.D.** *For any $T \in \mathrm{SL}(E)$,*

$$d_\angle(Tu, Tv) \le \|T^{-1}\|^2 \, d_\angle(u, v)$$

*for all $u, v \in E$.*

*Proof.*

$$\begin{aligned} d_\angle(Tu, Tv) &= \frac{|D(Tu, Tv)|}{\|Tu\|\|Tv\|} \\ &= \frac{\|u\|}{\|Tu\|} \, \frac{\|v\|}{\|Tv\|} \, \frac{|D(u,v)|}{\|u\|\|v\|} \\ &= \frac{\|u\|}{\|Tu\|} \, \frac{\|v\|}{\|Tv\|} \, d_\angle(u, v) \\ &\le \|T^{-1}\|^2 \, d_\angle(u, v). \end{aligned}$$

□

# E Standard punctures for translation surfaces

## E.1 Motivation

To see where the standard puncture shapes come from, we need to talk about the complex geometry of translation surfaces, whose only role until now has been a brief appearance in the proof of Proposition 3.5.H.

Every translation surface $\Sigma$ comes with a complex structure, which we get by identifying $\mathbb{R}^2$ with $\mathbb{C}$ in the usual way, and a complex-valued 1-form $\omega$, which sends horizontal unit vectors to 1 and vertical unit vectors to $i$. Observing that

$\omega = dz$ for any local translation chart $z\colon \Sigma \to \mathbb{C}$, we see that $\omega$ is holomorphic. Conversely, a complex 1-manifold equipped with a holomorphic 1-form is canonically a translation surface. Where $\omega$ has a zero of order $n$, the translation structure has a conical singularity of angle $2(n+1)\pi$.

The complex point of view suggests a natural class of translation surfaces that are non-compact, but still well-behaved. Putting a meromorphic 1-form on a compact Riemann surface defines a translation structure on the complement of the poles. The poles correspond to the ends of the translation surface, and the Riemann surface is its end compactification [27, §1]. The poles of a 1-form have a limited variety of behaviors, so the ends of the translation surface have a limited variety of shapes, which we'll call the *standard punctures*.

## E.2 First-order punctures

A first-order pole in $\omega$ makes a puncture shaped like a half-infinite cylinder. You can build one by rolling up a half-infinite rectangular strip and gluing its sides together:

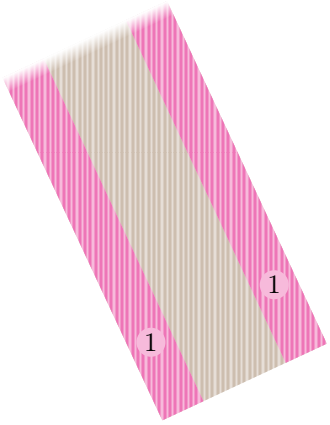


The translation structure of the cylinder you end up with is determined by two parameters: the width of the strip and its orientation in the plane. In most orientations, the vertical leaves spiral up or down the cylinder. When the strip is horizontal, the vertical leaves close up into circles.

## E.3 Higher-order punctures

The puncture created by a higher-order pole in $\omega$ can be glued together out of planes with quadrants cut away, like this:

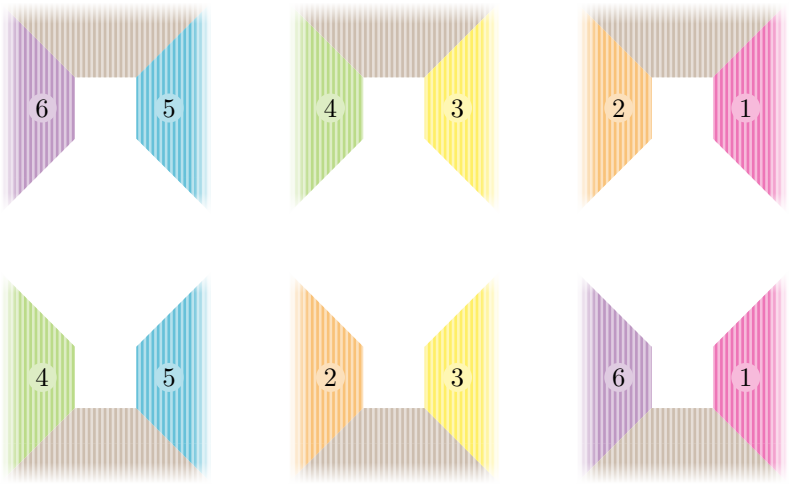


The notches at the centers of the pieces fit together into a polygonal hole that the rest of the surface can be connected to. For the pieces to fit together, the notches all have to be the same size, which can be adjusted to accommodate the part of the surface the puncture is supposed to hook up to.

Unlike first-order punctures, which come in a $\mathbb{C}^\times$-worth of shapes parameterized by the residues at their poles, higher-order punctures depend only on the orders of their poles.

## E.4 Counting ends

As I remarked earlier, a puncture is an end of the surface it lives in. Intuitively, you can think of it as a point on the boundary at infinity of the surface. The notion of an end is purely topological, however, and you might wonder if there's a notion of boundary at infinity that takes the translation structure into account. Here's one proposal, which is suitable at least for translation surfaces with punctures.

Let's say two vertical rays on a translation surface are *translation-equivalent* if they're connected by a continuous family of vertical rays.[11] A *vertical end* is an equivalence class of vertical rays. A generic first-order puncture, whose vertical leaves spiral up or down rather than closing up into circles, is a single vertical end. A higher-order puncture, on the other hand, comprises several vertical ends, one for each building block.

[11] To be formal about it, define a vertical ray to be a local isometry of $[0, \infty)$ into a vertical leaf. We can then say a continuous family of vertical rays is a continuous map from $[0, 1] \times [0, \infty)$ into the surface which restricts to a vertical ray when the first argument is fixed.

# References


[1] M. Mulase, "Geometry of character varieties of surface groups," *Research Institute for Mathematical Sciences Kokyuroku* (2008) , arXiv:0710.5263 [math.AG].

[2] J. Martínez and V. Muñoz, "E-polynomials of $\mathrm{SL}(2,\mathbb{C})$-character varieties of complex curves of genus 3," arXiv:1405.7120v2 [math.AG].

[3] V. Fock and A. Goncharov, "Moduli spaces of local systems and higher Teichmüller theory," arXiv:math/0311149 [math.AG].

[4] L. Hollands and A. Neitzke, "Spectral networks and Fenchel-Nielsen coordinates," arXiv:1312.2979 [math.GT].

[5] D. Gaiotto, G. Moore, and A. Neitzke, "Spectral networks," *Annales Henri Poincaré* **14** no. 7, (2013) 1643 – 1731.

[6] F. Bonahon and G. Dreyer, "Hitchin characters and geodesic laminations," arXiv:1410.0729v2 [math.GT].

[7] A. Furman, "On the multiplicative ergodic theorem for uniquely ergodic systems," *Annales de l'Institut Henri Poincaré (B)* **33** no. 6, (1997) 797 – 815.

[8] M. Viana, *Lectures on Lyapunov Exponents*, vol. 145 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 2014.

[9] D. Lenz, "Existence of non-uniform cocycles on uniquely ergodic systems," *Annales de l'Institut Henri Poincaré (B)* **40** no. 2, (2004) 197 – 206.

[10] The Stacks Project Authors, "*Stacks Project*." http://stacks.math.columbia.edu, 2013.

[11] H. Masur and S. Tabachnikov, "Rational billiards and flat structures," in *Handbook of Dynamical Systems*, B. Hasselblatt and A. Katok, eds., vol. 1A, ch. 13. 2002.

[12] J. Smillie, "The dynamics of billiard flows in rational polygons of dynamical systems," in *Dynamical Systems, Ergodic Theory and Applications*, Y. Sinai, ed., ch. 11. 2000.

[13] R. Gjerde and Ø. Johansen, "Bratteli-Vershik models for Cantor minimal systems associated to interval exchange transformations," *Mathematica Scandinavica* **90** no. 1, (2002) 87 – 100.

[14] G. Gruenhage, "MH 7750: Set theoretic topology." http://www.auburn.edu/~gruengf/fall14.html, 2014.

[15] M. Viana, "Lyapunov exponents of Teichmüller flows," in *Partially Hyperbolic Dynamics, Laminations, and Teichmüller Flow*, G. Forni, M. Lyubich, C. Pugh, and M. Shub, eds., vol. 51 of *Fields Institute Communications*. 2007.

[16] L. Marchese, "Khinchin theorem for interval exchange transformations," arXiv:1003.5883 [math.DS].

[17] L. Marchese, "Khinchin type condition for translation surfaces and asymptotic laws for the Teichmuller flow," arXiv:1003.5887 [math.DS].

[18] A. Zorich, "How do the leaves of a closed 1-form wind around a surface," in *Pseudoperiodic Topology*, V. Arnold, M. Kontsevich, and A. Zorich, eds., vol. 197 of *Translations of the AMS, Series 2*. 1999.

[19] D. Dumas and M. Wolf, "Projective structures, grafting, and measured laminations," arXiv:0712.0968 [math.DG].

[20] G. Dreyer, "Thurston's cataclysms for Anosov representations," arXiv:1301.6961 [math.GT].

[21] D. Gaiotto, G. Moore, and A. Neitzke, "Wall-crossing, Hitchin systems, and the WKB approximation," *Advances in Mathematics* **234** (2013) 239 – 403.

[22] V. Araujo, A. Bufetov, and S. Filip, "Hölder-continuity of Oseledets subspaces for the Kontsevich-Zorich cocycle," arXiv:1409.8167v2 [math].

[23] A. Fenyes, "Potentially cluster-like coordinates from dense spectral networks." https://www.ma.utexas.edu/users/afenyes/writing.html. Poster presented at Positive Grassmannians: Applications to integrable systems and super Yang-Mills scattering amplitudes. July 2015, CRM.

[24] S. Welstead, "Infinite products in a banach algebra," *Journal of Mathematical Analysis and Applications* **105** (1985) 523 – 532.

[25] K. D. Joshi, *Introduction to General Topology*. New Age International, 1983.

[26] E. Schechter, *Handbook of Analysis and its Foundations*. Academic Press, 1997.

[27] G. Peschke, "The theory of ends," *Nieuw Archief voor Wiskunde* **8** (1990) 1 – 12.